ipinfusion™

OcNOS®

Open Compute Network Operating System for Service Providers

Layer 3

Version 7.0.0

February 2026

IP Infusion Inc. Proprietary

©2026 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, CA 95054

+1 408-400-1900

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

Contents **3**

Preface **163**
About this Guide 163
Audience 163
Conventions 163
IP Infusion Product Release Version 163
Related Documentation 164
Feature Availability 164
Migration Guide 164
IP Maestro Support 164
Technical Support 164
Technical Sales 164
Technical Documentation 164
Documentation Disclaimer 165
Comments 165
Command Line Interface 166
Overview 166
Chapter Organization 166
Command Line Interface Help 167
Command Completion 168
Command Abbreviations 169
Command Line Errors 170
Command Negation 170
Syntax Conventions 171
Variable Placeholders 172
Command Description Format 173
Keyboard Operations 174
Show Command Modifiers 175
Begin Modifier 175
Include Modifier 176
Exclude Modifier 176
Redirect Modifier 177
Last Modifier 177
String Parameters 178
Command Modes 179
Command Mode Tree 180
Transaction-based Command-line Interface 181

Layer 3 Unicast Configuration **182**
Forwarding Plane Load Balancing 184
Enable Load Balancing 184
Topology 184

© 2026 IP Infusion Inc. Proprietary

Configuration 184
R1 - NSM 184
R1 - OSPF 185
R2 - OSPF 185
R3 - OSPF 185
Validation 186
R1 - OSPF 186
VLAN Interfaces 187
Overview 187
Topology 187
Create a VLAN Interface 188
Add IP Addresses to VLAN Interface 188
Display VLAN Interfaces 188
Layer 3 Link Aggregation 191
Overview 191
Topology 191
Configuration 191
Static Channel-group 193
IPv4Static Routes 196
Topology 196
Configuration 196
Validation 197
IPv6 Static Routs 201
Topology 201
Configuration 201
Validation 202
Static Route Discard Configuration 206
Overview 206
IPv4 Route Discard 206
Topology 206
Configuration 206
Validation 208
IPv6 Route Discard 209
Configuration 209
Validation 212
Static Route Object Tracking using IP SLA 214
Overview 214
Topology 214
Static Route Object Tracking IPv4 Configuration 214
RTR1 215
RTR2 216
RTR3 216
RTR4 216
Static Route Object Tracking IPv4 Validation 217
Static Route Object Tracking IPv6 Configuration 220
RTR1 220
RTR2 221

© 2026 IP Infusion Inc. Proprietary

RTR3 222
RTR4 222
Static Route Object Tracking IPv6 Validation 223
Route Monitor 224
Overview 224
Feature Characteristics 224
Benefits 224
Configuration 224
Prerequisites 224
Topology 225
IPv4 Route Monitor Configuration 225
IPv6 Route Monitor Configuration 228
Implementation Examples 232
Route Monitor Command 232
object-tracking 232
Troubleshooting 233
Glossary 233
Hybrid Switch Router Configuration 235
Overview 235
Routing and Switching 235
System Configuration 236
Hybrid Switch Router Possibilities 236
Configuring Layer 2 Interfaces 238
Configuring Layer 3 Interfaces 238
Policy Based Routing Configuration 240
Topology 240
Configurations - PBR IPv4 241
R1 241
R2 242
R3 242
R4 243
R5 243
Validation 244
Configuration - PBR IPv6 245
R1 245
R2 246
R3 247
R4 247
R5 248
Validation 249
PBR with Nexthop Tracking 250
IPv4 Configurations for PBR with Nexthop Tracking 250
R1 250
Validation 252
IPv6 Configurations for PBR with Nexthop Tracking 252
R1 252
Validation 254

TOS based Queue Distribution Configuration 255
Precedence-to-queue 255
Topology 255
Configuration 255
Validation 256
Precedence-to-precedence 256
Topology 256
Basic configuration 257
Validation 257
Internet Protocol SLA Configuration 258
Topology 258
Configuration 258
Configure IP Address 258
Configure IP SLA Configurations on PE 1 router 259
Validation 259
PE-1 259
EVPN AF Route-Maps and Route Filtering 261
Overview 261
Feature Characteristics 261
Benefits 261
Prerequisites 261
Limitation 262
Configuration 262
Validation 264
Implementation Examples 271
Commands 271
mac-list 272
match mac address list 273
match evpn-route-type 274
show mac-list 275
Glossary 275

Fundamental Layer 3 Command Reference 276

Fundamental Layer 3 Commands 279
automatic-router-id-selection enable 280
Command Syntax 280
Parameters 280
Default 280
Command Mode 280
Applicability 280
Examples 280
clear ip route 281
Command Syntax 281
Parameters 281
Command Mode 281
Applicability 281
Example 281

clear ip route kernel 282
Command Syntax 282
Parameters 282
Command Mode 282
Applicability 282
Example 282
clear ip route vrf 283
Command Syntax 283
Parameters 283
Command Mode 283
Applicability 283
Example 283
clear router-id 284
Command Syntax 284
Parameters 284
Command Mode 284
Applicability 284
Examples 284
debug rib 285
Command Syntax 285
Parameters 285
Disabled 286
Command Mode 286
Applicability 286
Example 286
description 287
Command Syntax 287
Parameters 287
Disabled 287
Command Mode 287
Applicability 287
Examples 287
fib retain 288
Command Syntax 288
Parameters 288
Default 288
Command Mode 288
Applicability 288
Examples 289
ip route 290
Command Syntax 290
Parameters 290
Default 291
Command Mode 291
Applicability 291
Examples 291
ip route vrf 293

Command Syntax 293
Parameters 293
Command Mode 293
Applicability 293
Example 293
ip urpf allow-default 294
Command Syntax 294
Parameter 294
Default 294
Command Mode 294
Applicability 294
Examples 294
ip urpf enable 295
Command Syntax 295
Parameter 295
Default 295
Command Mode 295
Applicability 295
Examples 295
ip verify unicast source reachable-via 296
Command Syntax 296
Parameters 296
Default 296
Command Mode 296
Applicability 296
Examples 296
ip vrf 297
Command Syntax 297
Parameter 297
Default 297
Command Mode 297
Applicability 297
Example 297
ipv6 route 298
Command Syntax 298
Parameters 298
Default 299
Command Mode 299
Applicability 299
Example 299
max-static-routes 300
Command Syntax 300
Parameters 300
Default 300
Command Mode 300
Applicability 300
Examples 300

maximum-paths ... 301
Command Syntax ... 301
Parameter ... 301
Default ... 301
Command Mode ... 301
Applicability ... 301
Example ... 301
router-id ... 302
Command Syntax ... 302
Parameters ... 302
Default ... 302
Command Mode ... 302
Applicability ... 302
Examples ... 302
show debugging rib ... 303
Command Syntax ... 303
Parameters ... 303
Command Mode ... 303
Applicability ... 303
Examples ... 303
show ip route track-table ... 304
Command Syntax ... 304
Parameters ... 304
Default ... 304
Command Mode ... 304
Applicability ... 304
Examples ... 304
show ip rpf ... 305
Command Syntax ... 305
Parameters ... 305
Default ... 305
Command Mode ... 305
Applicability ... 305
Example ... 305
show ipv6 route track-table ... 307
Command Syntax ... 307
Parameters ... 307
Default ... 307
Command Mode ... 307
Applicability ... 307
Examples ... 307
show ipv6 rpf ... 308
Command Syntax ... 308
Parameters ... 308
Default ... 308
Command Mode ... 308
Applicability ... 308

© 2026 IP Infusion Inc. Proprietary

Example 308
show rib txlist 310
Command Syntax 310
Parameters 310
Default 310
Command Mode 310
Applicability 310
Example 310
show router-id 312
Command Syntax 312
Parameters 312
Default 312
Command Mode 312
Applicability 312
Example 312
show running-config router 313
Command Syntax 313
Parameters 313
Default 313
Command Mode 313
Applicability 313
Example 313
show running-config router-id 314
Command Syntax 314
Parameters 314
Default 314
Command Mode 314
Applicability 314
Example 314
show running-config urpf 315
Command Syntax 315
Parameters 315
Default 315
Command Mode 315
Applicability 315
Example 315
show running-config vrf 316
Command Syntax 316
Parameters 316
Command Mode 316
Applicability 316
Example 316
snmp restart rib 317
Command Syntax 317
Parameters 317
Default 317
Command Mode 317

Applicability 317
Examples 317
IP Service Level Agreements Commands 318
clear ip sla statistics 319
Command Syntax 319
Parameters 319
Default 319
Command Mode 319
Applicability 319
Example 319
frequency 320
Command Syntax 320
Parameters 320
Default 320
Command Mode 320
Applicability 320
Example 320
icmp-echo 321
Command Syntax 321
Parameters 321
Default 321
Command Mode 321
Applicability 321
Example 321
ip sla 322
Command Syntax 322
Parameters 322
Default 322
Command Mode 322
Applicability 322
Example 322
ip sla schedule 323
Command Syntax 323
Parameters 323
Default 323
Command Mode 323
Applicability 323
Example 323
show ip sla statistics 324
Command Syntax 324
Parameters 324
Default 324
Command Mode 324
Applicability 324
Example 324
show ip sla summary 326
Command Syntax 326

Parameters 326
Default 326
Command Mode 326
Applicability 326
Example 326
show running-config ip sla 327
Command Syntax 327
Parameters 327
Default 327
Command Mode 327
Applicability 327
Example 327
threshold 328
Command Syntax 328
Parameters 328
Default 328
Command Mode 328
Applicability 328
Example 328
timeout 329
Command Syntax 329
Parameters 329
Default 329
Command Mode 329
Applicability 329
Example 329
Object Tracking Commands 330
delay up down 331
Command Syntax 331
Parameters 331
Default 331
Command Mode 331
Applicability 331
Example 331
object-tracking 332
Command Syntax 332
Parameters 332
Default 332
Command Mode 332
Applicability 332
Example 332
show running-config track 334
Command Syntax 334
Parameters 334
Default 334
Command Mode 334
Applicability 334

Example 334
show track 335
Command Syntax 335
Parameters 335
Default 335
Command Mode 335
Applicability 335
Example 335
show track <1-500> 336
Command Syntax 336
Parameters 336
Default 336
Command Mode 336
Applicability 336
Example 336
show track summary 337
Command Syntax 337
Parameters 337
Default 337
Command Mode 337
Applicability 337
Example 337
track ip sla reachability 338
Command Syntax 338
Parameters 338
Command Mode 338
Applicability 338
Example 338
Route-Map Commands 339
continue 341
Continue clause with match Clauses 341
Continue clause with set Clauses 341
Command Syntax 341
Parameter 341
Default 341
Command Mode 342
Applicability 342
Example 342
match as-path 343
Command Syntax 343
Parameter 343
Default 343
Command Mode 343
Applicability 343
Example 343
match community 344
Command Syntax 344

Parameters 344
Default 344
Command Mode 344
Applicability 344
Examples 344
match extcommunity 345
Command Syntax 345
Parameters 345
Default 345
Command Mode 345
Applicability 345
Example 345
match interface 346
Command Syntax 346
Parameter 346
Default 346
Command Mode 346
Applicability 346
Example 346
match ip address 347
Command Syntax 347
Parameters 347
Default 347
Command Mode 347
Applicability 347
Example 347
match ip address prefix-list 348
Command Syntax 348
Parameter 348
Default 348
Command Mode 348
Applicability 348
Examples 348
match ip next-hop 349
Command Syntax 349
Parameters 349
Default 349
Command Mode 349
Applicability 349
Examples 349
match ip next-hop prefix-list 350
Command Syntax 350
Parameter 350
Default 350
Command Mode 350
Applicability 350
Examples 350

© 2026 IP Infusion Inc. Proprietary

match ip peer 351
Command Syntax 351
Parameter 351
Default 351
Command Mode 351
Applicability 351
Examples 351
match ipv6 address 352
Command Syntax 352
Parameter 352
Default 352
Command Mode 352
Applicability 352
Examples 352
match ipv6 address prefix-list 353
Command Syntax 353
Parameter 353
Default 353
Command Mode 353
Applicability 353
Examples 353
match ipv6 next-hop 354
Command Syntax 354
Parameters 354
Default 354
Command Mode 354
Applicability 354
Examples 354
match ipv6 next-hop prefix-list 355
Command Syntax 355
Parameters 355
Default 355
Command Mode 355
Applicability 355
Examples 355
match ipv6 peer 356
Command Syntax 356
Parameter 356
Default 356
Command Mode 356
Applicability 356
Examples 356
match metric 357
Command Syntax 357
Parameters 357
Default 357
Command Mode 357

Applicability 357
Examples 357
match origin 358
Command Syntax 358
Parameters 358
Default 358
Command Mode 358
Applicability 358
Example 359
match route-map 359
Command Syntax 359
Parameters 359
Default 359
Command Mode 359
Applicability 359
Examples 359
match route-type 361
Command Syntax 361
Parameters 361
Default 361
Command Mode 361
Applicability 361
Examples 361
match tag 362
Command Syntax 362
Parameters 362
Default 362
Command Mode 362
Applicability 362
Examples 362
route-map 363
Command Syntax 363
Parameters 363
Default 363
Command Mode 363
Applicability 363
Example 363
set aggregator 364
Command Syntax 364
Parameters 364
Default 364
Command Mode 364
Applicability 364
Examples 364
set as-path 365
Command Syntax 365
Parameters 365

© 2026 IP Infusion Inc. Proprietary

Default 365
Command Mode 365
Applicability 365
Examples 365
set atomic-aggregate 366
Command Syntax 366
Parameters 366
Default 366
Command Mode 366
Applicability 366
Examples 366
set comm-list 367
Command Syntax 367
Parameters 367
Default 367
Command Mode 367
Applicability 367
Examples 367
set community 368
Command Syntax 368
Parameters 368
Default 368
Command Mode 368
Applicability 368
Examples 369
set dampening 370
Command Syntax 370
Parameters 370
Default 370
Command Mode 370
Applicability 370
Example 370
set extcommunity 371
Command Syntax 371
Parameters 371
Default 371
Command Mode 372
Applicability 372
Examples 372
set interface null0 373
Command Syntax 373
Parameters 373
Default 373
Command Mode 373
Applicability 373
Example 373
set ip next-hop 374

Command Syntax ... 374
Parameter ... 374
Default ... 374
Command Mode ... 374
Applicability ... 374
Examples ... 374
set ipv6 next-hop ... 376
Command Syntax ... 376
Parameters ... 376
Default ... 376
Command Mode ... 376
Applicability ... 376
Examples ... 376
set ip next-hop self ... 377
Command Syntax ... 377
Parameters ... 377
Command Mode ... 377
Applicability ... 377
Examples ... 377
set large-community ... 378
Command Syntax ... 378
Parameters ... 378
Default ... 378
Command Mode ... 378
Applicability ... 378
Examples ... 378
set level ... 380
Command Syntax ... 380
Parameters ... 380
Default ... 380
Command Mode ... 380
Applicability ... 380
Examples ... 380
set local-preference ... 381
Command Syntax ... 381
Default ... 381
Command Mode ... 381
Applicability ... 381
Example ... 381
set metric ... 382
Command Syntax ... 382
Parameters ... 382
Default ... 382
Command Mode ... 382
Applicability ... 382
Examples ... 382
set metric-type ... 383

© 2026 IP Infusion Inc. Proprietary

Command Syntax 383
Default 383
Command Mode 383
Applicability 383
Example 383
set origin 384
Command Syntax 384
Default 384
Command Mode 384
Applicability 384
Example 384
set originator-id 385
Command Syntax 385
Parameter 385
Default 385
Command Mode 385
Applicability 385
Example 385
set tag 386
Command Syntax 386
Parameter 386
Default 386
Command Mode 386
Applicability 386
Example 386
set vpnv4 next-hop 387
Command Syntax 387
Parameter 387
Default 387
Command Mode 387
Applicability 387
Example 387
set weight 388
Command Syntax 388
Parameter 388
Default 388
Command Mode 388
Applicability 388
Examples 388
show route-map 389
Command Syntax 389
Parameters 389
Command Mode 389
Applicability 389
Example 389
show running-config route-map 390
Command Syntax 390

Parameters 390
Command Mode 390
Applicability 390
Example 390
Source Interface Commands 391
ip source-interface 392
Command Syntax 392
Parameters 392
Default 393
Command Mode 393
Applicability 393
Example 393
ipv6 source-interface 394
Command Syntax 394
Parameters 394
Default 394
Command Mode 395
Applicability 395
Example 395
show ip source-interface detail 396
Command Syntax 396
Parameters 396
Default 396
Command Mode 396
Applicability 396
Example 396
show ipv6 source-interface detail 398
Command Syntax 398
Parameters 398
Default 398
Command Mode 398
Applicability 398
Example 398
show running-config ip source-interface 400
Command Syntax 400
Parameters 400
Default 400
Command Mode 400
Applicability 400
Example 400
show running-config ipv6 source-interface 401
Command Syntax 401
Parameters 401
Default 401
Command Mode 401
Applicability 401
Example 401

© 2026 IP Infusion Inc. Proprietary

Border Gateway Protocol Configuration 402
BGP 408
Enable BGP Routers in the Same Autonomous System 408
Topology 408
R1 408
R2 408
Validation 409
Enable BGP Between Different Autonomous Systems 409
Topology 410
Validation 410
Route-Map 411
Topology 412
Validation 413
Route-map Continue 413
Using Continue with Match Clauses 414
Using Continue with Set Actions 414
Topology 414
Configure Route-map continue on R1 414
Validation 417
Route-map Show Commands 419
Route-map Nesting 420
Topology 420
Configure Route-map 421
Validation 424
Route Reflector 424
Reduce the iBGP Mesh Inside an Autonomous System 424
Topology 424
Validation 426
Multiple Route Reflectors 428
Topology 429
Validation 432
BGP Confederations 438
Topology 438
Validation 440
Multiple Autonomous Systems 442
Outside Autonomous System 446
Dynamic BGP Peering 450
IPv4 IBGP Configuration 451
IPv4 IBGP VRF Configuration 455
IPv4 EBGP Configuration 459
IPv6 IBGP Configuration 462
IPV6 IBGP VRF Configuration 466
IPv6 EBGP Configuration 469
R2 471
VPNV4 Configuration 473
R2 475
R2 478

© 2026 IP Infusion Inc. Proprietary

R3 ... 478
Enable eBGP Multihop ... 479
Topology ... 479
Validation ... 480
Enable Peer Groups ... 482
Category A: Neighbor configuration allowed only at peer-group level ... 483
Category B: Neighbor configuration allowed at peer-group member level; precedence based ... 483
BGP Peer Groups for Address-Family IPv4 Unicast ... 484
R3 ... 490
Route Redistribution in BGP ... 493
Topology ... 494
Validation ... 494
Add Multiple Instances of the Same Autonomous System ... 494
Topology ... 495
Validation ... 497
Remove the Multi-Exit Disc Attribute from Update Messages ... 497
Topology ... 497
Removing Sent and Received MED values ... 498
Validation ... 499
BGP Four-Byte Autonomous System ... 500
4-Octet ASN Capability Enabled on R1 and R2 ... 500
Topology ... 500
4-Octet ASN Capability Enabled on R1 and Disabled on R2 ... 501
Topology ... 501
Topology ... 502
Validation ... 503
BGP Extended Community Attribute ... 504
Extended Community with a 2-Byte ASN ... 504
Extended Community with a 4-Byte ASN ... 513
Nexthop Tracking ... 521
Topology ... 521
Validation ... 523
Nexthop Tracking Delay Timer ... 523
Topology ... 524
Validation ... 526
BGP Distance ... 530
Topology ... 530
Validation ... 531
BGP Weight per Peer ... 533
Topology ... 533
Validation ... 534
OSPF as PE-CE Protocol for VPNs ... 536
Topology ... 536
Configuration ... 537
Validation ... 539
BGP Multipath for IPv4 ... 539

© 2026 IP Infusion Inc. Proprietary

Topology 540
Configuration 540
Validation 542
Multipath eBGP 543
Topology 543
Configuration 543
Validation 545
Multipath eiBGP 546
Topology 546
Configuration 546
Validation 549
BGP AS-PATH Multipath-relax 550
Topology 550
Configuration 551
Validation 552
BGP FIB Install (Selective Route Download) 553
Topology 554
Configuration 554
Validation 555
Route Target Constraint 559
Topology 559
Configuration 560
Validation 566
BGP Best Path Selection Process 569
BGP Dampening 570
Topology 571
IPv4 Configuration 571
Validation 573
IPv6 Configuration 574
Validation 575
BGP Authentication 576
Topology 577
Configuration 577
Validation 578
BGP Unnumbered 579
Overview 579
Topology 580
Configuration 580
Validation 583
BGP Blackhole Community Attribute 587
Topology 587
Configuration 587
Validation 589
Accumulated Interior Gateway Protocol 590
Overview 590
Topology 590
Validation 596

© 2026 IP Infusion Inc. Proprietary

Extended Community Attribute ... 597
Validation ... 600
BGP Graceful Restart Configuration ... 601
Topology ... 601
Validation ... 603
Validation After BGP Graceful Restart ... 606
BGP Automated Dynamic Route Policy Update ... 608
Overview ... 608
Benefits ... 608
Prerequisites ... 609
Configuring BGP auto-policy-soft-reset ... 609
Topology ... 609
Procedure ... 609
Validation ... 613
Implementation Examples ... 617
CLI Commands ... 618
bgp auto-policy-soft-reset enable ... 619
debug bgp soft-clear ... 620
Glossary ... 620
BGP4+ ... 621
Enable iBGP Peering Using a Global Address ... 621
Topology ... 621
Configuration ... 621
R1 ... 621
R2 ... 622
Validation ... 622
R1 ... 622
R2 ... 623
Enable iBGP Peering Using Link-local Address ... 624
Topology ... 624
Configuration ... 624
R1 ... 624
R2 ... 625
Validation ... 625
R1 ... 625
R2 ... 626
Enable eBGP Peering Between Different Autonomous Systems ... 626
Topology ... 627
Configuration ... 627
R1 ... 627
R2 ... 627
Validation ... 628
R1 ... 628
R2 ... 628
Route-Map ... 629
Topology ... 629
Configuration ... 630

R1 630
R2 631
Validation 631
R1 631
R2 632
Route Reflector 633
Topology 633
Configuration 633
RR (R5) 633
RR Client 1 (R2) 634
RR Client 2 (R4) 635
Validation 635
Confederations 637
Topology 637
Configuration 637
R2 637
R5 638
R7 639
Validation 640
BGP4+ Graceful Restart 642
Topology 642
Configuration 642
R1 642
R2 643
Validation 643
R1 643
R2 645
Performing BGP Graceful Restart 646
Restarting the BGP Daemon 647
Configure BGP4+ Distance 647
Topology 648
Configuration 648
R1 648
R2 648
R3 649
Validation 649
BGP4+ Graceful Reset 652
Topology 652
Configuration 652
R1 652
R2 652
Validation 653
VPNv4 NLRI with IPv6 Nexthop 656
Topology 656
Configuration 656
R1 656
R2 657

© 2026 IP Infusion Inc. Proprietary

Validation 659
R1 659
R2 660
BGP Labeled Unicast 662
BGP Labeled Unicast as Transport Configuration 663
Topology 663
PE1 Configuration 663
ABR1 Configuration 665
P1 Configuration 666
ABR2 Configuration 667
PE2 Configuration 669
BGP LU Transport Validation 671
BGP Labeled Unicast with Seamless MPLS Configuration 675
Topology 675
PE1 Configuration 675
P1 Configuration 677
ABR1 Configuration 678
P2 Configuration 680
ABR2 Configuration 681
P3 Configuration 683
PE2 Configuration 684
BGP LU with Seamless MPLS Validation 686
BGP Labeled Unicast with Inter-AS Configuration 695
Topology 695
PE1 Configuration 695
P1 Configuration 697
ABR1 Configuration 698
ABR2 Configuration 700
P2 Configuration 701
PE2 Configuration 702
BGP LU with Inter-AS Validation 704
BGP Peer Groups for Address-Family IPv4 Labeled Unicast 709
Topology 709
R1 Configuration 709
R2 Configuration 710
R3 Configuration 710
BGP Peer Groups for Address-Family IPv4 LU Validation 711
Peer Group Outbound Policies Inheritance 713
Peer Group Inbound Policies Inheritance 716
BGP Labeled Unicast with Implicit Null Label for Local Routes 720
Overview 720
Benefits 720
Prerequisites 720
Topology 721
Configuration 721
BGP Labeled Unicast as Transport with Implicit Null Label 721
Sample show running-config Output 721

© 2026 IP Infusion Inc. Proprietary

Validation 723
PE1 723
ABR1 725
P1 725
PE2 726
Glossary 728
BGP MPLS Next Hop Tracking Configuration 729
Topology 729
Configuration 729
PE1 729
ABR1 731
ABR2 733
PE2 735
Validation 737
PE1 737
PE2 738
ABR1 739
ABR2 740
BGP Labeled Unicast Next Hop Self in Route-Map 742
Overview 742
Feature Characteristics 742
Limitations 742
Prerequisites 742
Configuration 742
Configuration Snapshot 744
Commands 802
set ip next-hop self 802
BGP Additional Paths for VPNv4 803
Overview 803
Normal BGP Behavior 803
BGP Behavior with ADDPATH 803
Topology 804
Initial Configuration 804
CE1 Configuration 804
PE1Configuration 805
RR Configuration 806
PE2 Configuration 808
CE2 Configuration 810
Additional Paths Capability Send/Receive/Send-Receive and Additional Paths Selection All/Best 2/Best 3 811
Additional Paths for 6PE 821
RR 831
PE1 832
PE2 832
PE3 832
PE4 833
RR 833

IPV6 LABELED-UNICAST Prefix Advertised to All 6pe Peers With All Addpaths: 836
PE1 836
PE2 837
PE3 837
PE4 838
CE1 839
RR 839
Validation 839
RR 839
PE1 840
PE2 841
PE3 841
PE4 842
RR 843
Validation 843
RR 843
PE1 844
PE2 845
PE3 845
PE4 846
RR 846
Validation 847
PE1 847
RR 847
Validation 848
PE1 848
RR 848
Validation 849
PE1 849
RR 850
Validation 850
PE1 850
CE1 850
CE2 851
PE1 851
PE2 852
RR 852
PE3 854
CE3 854
Additional Paths for 6VPE 854
Topology 855
Configuration 855
BGP Labeled Unicast IPv4 with Add Path 887
Overview 887
Topology 887
Interface Connections 887
Configurations 888

Validation for BGP Labeled Unicast with Additional Paths 918
BGP4+ Additional Paths Configuration 971
Overview 971
Normal BGP Behavior 971
BGP Behavior with Additional Paths 971
Topology 973
R1 973
R2 974
R3 975
R4 976
Additional Paths at the Global Level 978
R2 978
R3 978
Additional Paths at the Neighbor Level 979
R2 979
R3 979
Validation 980
Additional Paths Send and Receive at Address-family level 984
R2 984
R3 984
Additional Paths at the Neighbor Level 985
R2 985
R3 985
Validation 986
Selection of all Additional Paths at the Address-family Level 991
R2 991
Selection of all Additional Paths at the Neighbor Level 992
R2 992
Validation 992
Selection of Best 2 Additional Paths at AF Level 994
R2 994
Selection of Best 2 Additional Paths at the Neighbor Level 995
R2 995
Validation 995
Selection of Best 3 Additional Paths at the AF Level 997
R2 997
Selection of Best 3 Additional Paths at the Neighbor Level 998
R2 998
Validation 998
BGP Flowspec Support for IPv4 1000
BGP Flowspec IPv4 Configuration 1000
Topology 1000
CE1 Configuration 1000
PE1 Configuration 1002
P1 Configuration 1005
P2 Configuration 1006
PE2 Configuration 1007

© 2026 IP Infusion Inc. Proprietary

CE2 Configuration 1010
BGP Flowspec IPv4 Validation 1012
BGP Flowspec Configuration With RR 1060
Topology 1060
CE1 Configuration 1060
PE1 Configuration 1062
RR2 Configuration 1066
RR1 Configuration 1068
PE2 Configuration 1070
CE2 Configuration 1073
BGP Flowspec Configuration With RR Validation 1075
TCP MSS configuration for BGP neighbors 1128
Overview 1128
Feature Characteristics 1128
Benefits 1128
Prerequisites 1129
Configuration 1129
Topology 1129
Configuration 1129
Validation 1131
PE1 1131
RR1 1132
New CLI Commands 1132
neighbor tcp-mss 1132
Abbreviations 1133
Glossary 1134
BGP ORF Prefix-List VPNV4 Address 1135
Overview 1135
Feature Characteristics 1135
Configuration 1135
Topology 1135
PE1 1136
P 1137
RR 1138
PE2 1139
Validation 1140
Glossary 1146
BGP Bogon Prefix Filtering IPv4 1148
Overview 1148
Feature Characteristics 1148
Benefits 1148
BGP Bogon Prefix Filtering Configuration 1148
Topology 1149
Prerequisites 1149
Configure Invalid IP Addresses on R1 and R2 1150
Enable BGP Bogon Filtering on R1 and R2 1151
Validation 1151

BGP Bogon Prefix Filtering Command 1151
bgp enable-bogon-filtering 1151
Glossary 1152
BGP Large Communities 1153
Overview 1153
Feature Characteristics 1153
Benefits 1153
BGP Large Communities Configuration 1153
Topology 1154
Prerequisites 1154
BGP Large Communities Routers Configurations 1159
Running Configurations 1165
Implementation Example 1171
BGP Large Communities Commands 1171
set large-community 1172
ip large-community-list 1173
match large-community 1175
set large-comm-list delete 1177
show ip large-community-list 1179
Troubleshooting 1180
Glossary 1180
Hide the Remote AS using the neighbor local-as Command 1181
Overview 1181
Feature Characteristics 1181
Benefits 1181
Configuration 1181
Topology 1181
Validation 1182
Revised CLI Commands 1184
neighbor local-as 1184
Abbreviations 1185
BGP Aggregate-Address in VRF Address-family 1186
Overview 1186
Benefits 1186
Prerequisites 1186
Configuration for Aggregate-Address in VRF Address-Family with IPv4 1186
Topology 1186
PE1 1187
PE2 1188
CE1 1189
Running Configuration with IPv4 1190
PE1 1190
PE2 1191
CE1 1192
Validation 1192
Configuration for Aggregate-Address in VRF Address-Family with IPv6 1193
Topology 1193

PE1 1193
PE2 1194
Running Configuration with IPv6 1195
PE1 1195
PE2 1196
Validation 1196
BGP RPKI-Based Route Validation 1198
Overview 1198
Feature Characteristics 1198
Benefits 1198
Configuration 1199
Prerequisites 1199
Topology 1199
Validation 1201
Implementation Examples 1203
Commands 1203
bgp rpki server 1204
bgp origin-as validation-enable 1206
bgp origin-as bestpath 1207
match rpki 1208
show bgp rpki table ipv4 1209
show bgp rpki table ipv6 1211
show bgp origin-as validity ipv4 1213
show bgp origin-as validity ipv6 1214
show bgp rpki server 1216
Glossary 1217
Recursive Next-hop Resolution with MPLS Next-hop 1219
Overview 1219
Feature Characteristics 1219
Benefits 1219
Limitation 1219
Configuration 1219
Configuration with MPLS LDP 1219
Configuration with Segment Routing MPLS (SR-MPLS) 1224
On PE1 1229
On PE2 1230
From PE1 to CE2 1231
From PE2 to CE1 1231
Implementation Example 1232
Use Case: Simplified MPLS Core for BGP-Free Operation 1232
Revised Commands 1232
show ip route mpls 1232
show ip route database 1232
show rib txlist 1232
Troubleshooting 1233
show rib txlist 1234
Glossary 1235

© 2026 IP Infusion Inc. Proprietary

Border Gateway Protocol Command Reference 1237
BGP Commands 1244
address-family ipv4 1248
Command Syntax 1248
Parameters 1248
Applicability 1249
Command Mode 1249
Applicability 1249
Examples 1249
aggregate-address 1250
Command Syntax 1250
Parameters 1250
Default 1250
Command Mode 1250
Applicability 1250
Examples 1250
auto-summary 1252
Command Syntax 1252
Parameters 1252
Default 1252
Command Mode 1252
Applicability 1252
Examples 1252
bgp aggregate-nexthop-check 1253
Command Syntax 1253
Parameters 1253
Default 1253
Command Mode 1253
Applicability 1253
Examples 1253
bgp always-compare-med 1254
Command Syntax 1254
Parameters 1254
Default 1254
Command Mode 1254
Applicability 1254
Examples 1254
bgp as-local-count 1255
Command Syntax 1255
Parameter 1255
Default 1255
Command Mode 1255
Applicability 1255
Examples 1255
bgp bestpath as-path ignore 1256
Command Syntax 1256
Parameters 1256

© 2026 IP Infusion Inc. Proprietary

Default 1256
Command Mode 1256
Applicability 1256
Examples 1256
bgp bestpath as-path multipath-relax 1257
Command Syntax 1257
Parameters 1257
Default 1257
Command Mode 1257
Applicability 1257
Examples 1257
bgp bestpath compare-confed-aspath 1258
Command Syntax 1258
Parameters 1258
Default 1258
Command Mode 1258
Applicability 1258
Examples 1258
bgp bestpath compare-routerid 1259
Command Syntax 1259
Parameters 1259
Default 1259
Command Mode 1259
Applicability 1259
Examples 1259
bgp bestpath dont-compare-originator-id 1260
Command Syntax 1260
Parameters 1260
Default 1260
Command Mode 1260
Applicability 1260
Examples 1260
bgp bestpath med 1261
Command Syntax 1261
Parameters 1261
Command Mode 1261
Default 1261
Applicability 1261
Examples 1262
bgp bestpath tie-break-on-age 1263
Command Syntax 1263
Parameters 1263
Default 1263
Command Mode 1263
Applicability 1263
Examples 1263
bgp client-to-client reflection 1264

Command Syntax 1264
Parameters 1264
Default 1264
Command Mode 1264
Applicability 1264
Examples 1264
bgp cluster-id 1265
Command Syntax 1265
Parameters 1265
Default 1265
Command Mode 1265
Applicability 1265
Examples 1265
bgp confederation vrf identifier 1266
Command Syntax 1266
Parameter 1266
Default 1266
Command Mode 1266
Applicability 1266
Examples 1266
bgp confederation vrf peers 1267
Command Syntax 1267
Parameter 1267
Default 1267
Command Mode 1267
Applicability 1267
Examples 1267
bgp dampening 1268
Command Syntax 1268
Parameters 1268
Defaults 1268
Command Mode 1268
Applicability 1269
Examples 1269
bgp default local-preference 1270
Command Syntax 1270
Parameter 1270
Default 1270
Command Mode 1270
Applicability 1270
Examples 1270
bgp deterministic-med 1271
Command Syntax 1271
Parameters 1271
Default 1271
Command Mode 1271
Applicability 1271

Examples 1271
bgp enforce-first-as 1273
Command Syntax 1273
Parameters 1273
Default 1273
Command Mode 1273
Applicability 1273
Examples 1273
bgp extended-asn-cap 1274
Command Syntax 1274
Parameters 1274
Default 1274
Command Mode 1274
Applicability 1274
Examples 1274
bgp fast-external-failover 1275
Command Syntax 1275
Parameters 1275
Default 1275
Command Mode 1275
Applicability 1275
Example 1275
bgp implicit-null 1276
Command Syntax 1276
Parameters 1276
Default 1276
Command Mode 1276
Applicability 1276
Example 1276
bgp log-neighbor-changes 1278
Command Syntax 1278
Parameters 1278
Default 1278
Command Mode 1278
Applicability 1279
Example 1279
mpls-nexthop-tracking 1280
Command Syntax 1280
Parameters 1280
Command Mode 1280
Applicability 1280
Examples 1280
bgp nexthop-trigger delay 1281
Command Syntax 1281
Parameter 1281
Default 1281
Command Mode 1281

© 2026 IP Infusion Inc. Proprietary

Applicability ... 1281
Examples ... 1281
bgp nexthop-trigger enable ... 1282
Command Syntax ... 1282
Parameters ... 1282
Default ... 1282
Command Mode ... 1282
Applicability ... 1282
Examples ... 1282
bgp rfc1771-path-select ... 1283
Command Syntax ... 1283
Parameters ... 1283
Default ... 1283
Command Mode ... 1283
Applicability ... 1283
Examples ... 1283
bgp rfc1771-strict ... 1284
Command Syntax ... 1284
Parameters ... 1284
Default ... 1284
Command Mode ... 1284
Applicability ... 1284
Examples ... 1284
bgp router-id ... 1285
Command Syntax ... 1285
Parameter ... 1285
Default ... 1285
Command Mode ... 1285
Applicability ... 1285
Examples ... 1285
bgp scan-time ... 1286
Command Syntax ... 1286
Parameter ... 1286
Default ... 1286
Command Mode ... 1286
Applicability ... 1286
Examples ... 1286
table-map ... 1287
Command Syntax ... 1287
Parameter ... 1287
Default ... 1287
Command Mode ... 1287
Applicability ... 1287
Examples ... 1287
bgp update-delay ... 1289
Command Syntax ... 1289
Parameters ... 1289

Default ... 1289
Command Mode ... 1289
Applicability ... 1289
Example ... 1289
bgp unnumbered-mode ... 1290
Command Syntax ... 1290
Parameters ... 1290
Command Mode ... 1290
Applicability ... 1290
Examples ... 1290
clear bgp (A.B.C.D|X:X::X:X) ... 1291
Command Syntax ... 1291
Parameters ... 1291
Command Mode ... 1291
Applicability ... 1291
Example ... 1291
clear bgp * ... 1292
Command Syntax ... 1292
Parameters ... 1292
Command Mode ... 1293
Applicability ... 1293
Examples ... 1293
clear bgp * l2vpn evpn ... 1294
Command Syntax ... 1294
Parameters ... 1294
Default ... 1294
Command Mode ... 1294
Applicability ... 1294
Examples ... 1294
clear bgp <1-4294967295> ... 1295
Command Syntax ... 1295
Parameters ... 1295
Command Mode ... 1296
Applicability ... 1296
Examples ... 1296
clear bgp <1-4294967295> l2vpn evpn ... 1297
Command Syntax ... 1297
Parameters ... 1297
Default ... 1297
Command Mode ... 1297
Applicability ... 1297
Examples ... 1297
clear bgp A.B.C.D l2vpn evpn ... 1298
Command Syntax ... 1298
Parameters ... 1298
Default ... 1298
Command Mode ... 1298

Applicability 1298
Examples 1298
clear bgp dampening 1299
Command Syntax 1299
Parameters 1299
Command Mode 1299
Applicability 1299
Examples 1299
clear bgp external 1300
Command Syntax 1300
Parameters 1300
Command Mode 1301
Applicability 1301
Example 1301
clear bgp flap-statistics 1302
Command Syntax 1302
Parameters 1302
Command Mode 1302
Applicability 1302
Examples 1302
clear bgp peer-group 1303
Command Syntax 1303
Parameters 1303
Command Mode 1304
Applicability 1304
Examples 1304
clear bgp peer-group WORD l2vpn evpn 1305
Command Syntax 1305
Parameters 1305
Default 1305
Command Mode 1305
Applicability 1305
Examples 1305
clear bgp statistics 1306
Command Syntax 1306
Parameters 1306
Command Mode 1306
Applicability 1306
Examples 1306
clear ip bgp * (description LINE|)* 1307
Command Syntax 1307
Parameters 1307
Command Mode 1307
Applicability 1307
Examples 1307
clear ip bgp (A.B.C.D|X:X::X:X|WORD) (description LINE|) 1308
Command Syntax 1308

Parameters 1308
Command Mode 1308
Applicability 1308
Examples 1308
clear ip bgp A.B.C.D 1309
Command Syntax 1309
Parameters 1309
Command Mode 1310
Applicability 1310
Examples 1310
clear ip bgp A.B.C.D vrf 1311
Command Syntax 1311
Parameters 1311
Command Mode 1311
Applicability 1311
Examples 1311
clear ip bgp all vrf (VRFNAME|all|default) (description LINE|) 1312
Command Syntax 1312
Parameters 1312
Command Mode 1312
Applicability 1312
Examples 1312
clear ip bgp table-map 1313
Command Syntax 1313
Parameters 1313
Command Mode 1313
Applicability 1313
Examples 1313
clear ip bgp vrf WORD (A.B.C.D|X:X::X:X|WORD) (description LINE|) 1314
Command Syntax 1314
Parameters 1314
Command Mode 1314
Applicability 1314
Examples 1314
debug bgp 1315
Command Syntax 1315
Parameters 1315
Command Mode 1316
Applicability 1316
Examples 1316
distance 1317
Command Syntax 1317
Parameters 1317
Default 1317
Command Mode 1317
Applicability 1317
Example 1318

exit-address-family 1319
Command Syntax 1319
Parameters 1319
Default 1319
Command Mode 1319
Applicability 1319
Examples 1319
ip as-path access-list 1320
Command Syntax 1320
Parameters 1320
Default 1320
Command Mode 1320
Applicability 1320
Examples 1320
ip community-list <1-99> 1321
Command Syntax 1321
Parameters 1321
Default 1321
Command Mode 1321
Applicability 1321
Examples 1321
ip community-list <100-500> 1322
Command Syntax 1322
Parameters 1322
Default 1322
Command Mode 1322
Applicability 1322
Examples 1322
ip community-list expanded 1323
Command Syntax 1323
Parameters 1323
Default 1323
Command Mode 1323
Applicability 1323
Examples 1323
ip community-list standard 1324
Command Syntax 1324
Parameters 1324
Default 1324
Command Mode 1324
Applicability 1325
Examples 1325
ip community-list WORD 1326
Command Syntax 1326
Parameters 1326
Default 1326
Command Mode 1326

Applicability 1326
Examples 1327
ip extcommunity-list <1-99> 1328
Command Syntax 1328
Parameters 1328
Default 1328
Command Mode 1328
Applicability 1328
Examples 1328
ip extcommunity-list <100-500> 1329
Command Syntax 1329
Parameters 1329
Default 1329
Command Mode 1329
Applicability 1329
Examples 1329
ip extcommunity-list expanded 1330
Command Syntax 1330
Parameters 1330
Default 1330
Command Mode 1330
Applicability 1330
Examples 1330
ip extcommunity-list standard 1331
Command Syntax 1331
Parameters 1331
Default 1331
Command Mode 1331
Applicability 1331
Examples 1331
l2vpn-unnumbered-mode 1333
Command Syntax 1333
Parameters 1333
Applicability 1333
Command Mode 1333
Applicability 1333
Examples 1333
match ip peer 1334
Command Syntax 1334
Parameters 1334
Default 1334
Command Mode 1334
Applicability 1334
Examples 1334
max-paths 1335
Command Syntax 1335
Parameters 1335

Default 1335
Command Mode 1335
Applicability 1335
Examples 1335
max-paths eigbp 1336
Command Syntax 1336
Parameters 1336
Default 1336
Command Mode 1336
Applicability 1336
Example 1336
neighbor activate 1338
Command Syntax 1338
Parameters 1338
Default 1338
Command Mode 1338
Applicability 1338
Examples 1338
neighbor advertisement-interval 1339
Command Syntax 1339
Parameters 1339
Default 1339
Command Mode 1339
Applicability 1339
Example 1339
neighbor allowas-in 1340
Command Syntax 1340
Parameters 1340
Default 1340
Command Mode 1341
Applicability 1341
Example 1341
neighbor as-origination-interval 1342
Command Syntax 1342
Parameters 1342
Default 1342
Command Mode 1342
Applicability 1342
Example 1342
neighbor attribute-unchanged 1343
Command Syntax 1343
Parameters 1343
Default 1343
Command Mode 1343
Applicability 1343
Example 1343
neighbor authentication-key 1345

Command Syntax 1345
Parameters 1345
Default 1345
Command Mode 1345
Applicability 1345
Example 1345
neighbor capability graceful-restart 1347
Command Syntax 1347
Parameters 1347
Default 1347
Command Mode 1347
Applicability 1347
Examples 1347
neighbor capability orf prefix-list 1348
Command Syntax 1348
Parameters 1348
Default 1348
Command Mode 1348
Applicability 1349
Examples 1349
neighbor collide-established 1350
Command Syntax 1350
Parameters 1350
Default 1350
Command Mode 1350
Applicability 1350
Example 1350
neighbor de-activate 1352
Command Syntax 1352
Parameters 1352
Default 1352
Command Mode 1352
Applicability 1352
Examples 1352
neighbor default-originate 1354
Command Syntax 1354
Parameters 1354
Default 1354
Command Mode 1354
Applicability 1354
Examples 1354
neighbor description 1356
Command Syntax 1356
Parameters 1356
Default 1356
Command Mode 1356
Applicability 1356

© 2026 IP Infusion Inc. Proprietary

Examples 1356
neighbor disallow-infinite-holdtime 1358
Command Syntax 1358
Parameters 1358
Default 1358
Command Mode 1358
Applicability 1358
Examples 1358
neighbor distribute-list 1359
Command Syntax 1359
Parameters 1359
Default 1359
Command Mode 1359
Applicability 1359
Examples 1360
neighbor ebgp-multihop 1361
Command Syntax 1361
Parameters 1361
Default 1361
Command Mode 1361
Applicability 1361
Examples 1361
neighbor enforce-multihop 1363
Command Syntax 1363
Parameters 1363
Default 1363
Command Mode 1363
Applicability 1363
Examples 1363
neighbor extended-optional-param 1364
Command Syntax 1364
Parameters 1364
Default 1364
Command Mode 1364
Applicability 1364
Example 1364
neighbor fall-over bfd 1365
Command Syntax 1365
Parameters 1365
Default 1365
Command Mode 1365
Applicability 1365
Example 1365
neighbor filter-list 1366
Command Syntax 1366
Parameters 1366
Default 1366

Command Mode 1366
Applicability 1366
Examples 1366
neighbor g-shut 1368
Command Syntax 1368
Parameters 1368
Default 1368
Command Mode 1368
Applicability 1368
Examples 1368
neighbor g-shut-timer 1369
Command Syntax 1369
Parameters 1369
Default 1369
Command Mode 1369
Applicability 1369
Examples 1369
neighbor limit 1370
Command Syntax 1370
Parameters 1370
Default 1370
Command Mode 1370
Applicability 1370
Example 1370
neighbor local-as 1371
Command Syntax 1371
Parameters 1371
Default 1371
Command Mode 1371
Applicability 1372
Example 1372
neighbor maximum-prefix 1373
Command Syntax 1373
Parameters 1373
Default 1373
Command Mode 1373
Applicability 1373
Examples 1374
neighbor next-hop-self 1375
Command Syntax 1375
Parameters 1375
Default 1375
Command Mode 1375
Applicability 1375
Examples 1375
neighbor optional-as 1376
Command Syntax 1376

Parameters 1376
Default 1376
Command Mode 1376
Applicability 1376
Examples 1376
neighbor override-capability 1377
Command Syntax 1377
Parameters 1377
Default 1377
Command Mode 1377
Applicability 1377
Examples 1377
neighbor passive 1378
Command Syntax 1378
Parameters 1378
Default 1378
Command Mode 1378
Applicability 1378
Examples 1378
neighbor peer-group 1379
Command Syntax 1379
Parameters 1379
Default 1379
Command Mode 1379
Applicability 1379
Example 1380
neighbor WORD peer-group range 1383
Command Syntax 1383
Parameters 1383
Default 1383
Command Mode 1383
Applicability 1383
Example 1383
neighbor port 1384
Command Syntax 1384
Parameters 1384
Default 1384
Command Mode 1384
Applicability 1384
Examples 1384
neighbor prefix-list 1385
Command Syntax 1385
Parameters 1385
Default 1385
Command Mode 1385
Applicability 1385
Examples 1386

neighbor remote-as 1387
Command Syntax 1387
Parameters 1387
Default 1387
Command Mode 1387
Applicability 1387
Example 1388
neighbor remove-private-AS 1389
Command Syntax 1389
Parameters 1389
Default 1389
Command Mode 1389
Applicability 1389
Example 1389
neighbor restart-time 1390
Command Syntax 1390
Parameters 1390
Default 1390
Command Mode 1390
Applicability 1390
Example 1390
neighbor route-map 1391
Command Syntax 1391
Parameters 1391
Default 1391
Command Mode 1391
Applicability 1392
Examples 1392
neighbor route-reflector-client 1393
Command Syntax 1393
Parameters 1393
Default 1393
Command Mode 1393
Applicability 1393
Examples 1393
neighbor route-server-client 1395
Command Syntax 1395
Parameters 1395
Default 1395
Command Mode 1395
Applicability 1395
Example 1395
neighbor send-community 1396
Command Syntax 1396
Parameters 1396
Default 1396
Command Mode 1396

© 2026 IP Infusion Inc. Proprietary

Applicability 1396
Examples 1397
neighbor send-community large 1398
Command Syntax 1398
Parameters 1398
Default 1398
Command Mode 1398
Applicability 1398
Examples 1398
neighbor shutdown 1399
Command Syntax 1399
Parameters 1399
Default 1399
Command Mode 1399
Applicability 1399
Examples 1399
neighbor strict-capability-match 1401
Command Syntax 1401
Parameters 1401
Default 1401
Command Mode 1401
Applicability 1401
Examples 1401
neighbor timers 1402
Command Syntax 1402
Parameters 1402
Defaults 1402
Command Mode 1402
Applicability 1403
Examples 1403
neighbor unsuppress-map 1404
Command Syntax 1404
Parameters 1404
Default 1404
Command Mode 1404
Applicability 1404
Example 1404
neighbor update-source 1406
Command Syntax 1406
Parameters 1406
Default 1406
Command Mode 1406
Applicability 1406
Examples 1406
neighbor version 1407
Command Syntax 1407
Parameters 1407

Default 1407
Command Mode 1407
Applicability 1407
Examples 1407
neighbor weight 1409
Command Syntax 1409
Parameters 1409
Default 1409
Command Mode 1409
Applicability 1409
Examples 1410
neighbor WORD peer-group 1411
Command Syntax 1411
Parameters 1411
Default 1411
Command Mode 1411
Applicability 1411
Example 1411
network 1412
Command Syntax 1412
Parameters 1412
Default 1412
Command Mode 1413
Applicability 1413
Examples 1413
network synchronization 1414
Command Syntax 1414
Parameters 1414
Default 1414
Command Mode 1414
Applicability 1414
Examples 1414
redistribute 1415
Command Syntax 1415
Parameters 1415
Default 1415
Command Mode 1415
Applicability 1415
Examples 1416
restart bgp graceful 1417
Command Syntax 1417
Parameters 1417
Default 1417
Command Mode 1417
Applicability 1417
Examples 1417
router bgp 1418

Command Syntax 1418
Parameters 1418
Default 1418
Command Mode 1418
Applicability 1418
Examples 1418
snmp restart bgp 1419
Command Syntax 1419
Parameters 1419
Default 1419
Command Mode 1419
Applicability 1419
Examples 1419
synchronization 1420
Command Syntax 1420
Parameters 1420
Default 1420
Command Mode 1420
Applicability 1420
Examples 1420
timers bgp 1421
Command Syntax 1421
Parameters 1421
Default 1421
Command Mode 1421
Applicability 1421
Examples 1421
v4-unnumbered-mode 1422
Command Syntax 1422
Parameters 1422
Applicability 1422
Command Mode 1422
Applicability 1422
Examples 1422
BGP4+ Commands 1423
address-family ipv6 1424
Command Syntax 1424
Parameters 1424
Applicability 1424
Command Mode 1424
Applicability 1424
Examples 1425
aggregate-address X:X::X:X/M 1426
Command Syntax 1426
Parameters 1426
Command Mode 1426
Default 1426

Applicability 1426
Examples 1426
clear bgp * ipv6 1428
Command Syntax 1428
Parameters 1428
Command Mode 1429
Applicability 1429
Examples 1429
clear bgp ipv6 (A.B.C.D|X:X::X:X) 1430
Command Syntax 1430
Parameters 1430
Command Mode 1430
Applicability 1430
Example 1430
clear bgp ipv6 <1-4294967295> 1431
Command Syntax 1431
Parameters 1431
Command Mode 1431
Applicability 1431
Example 1431
clear bgp ipv6 external 1432
Command Syntax 1432
Parameters 1432
Command Mode 1432
Applicability 1432
Example 1432
clear bgp ipv6 peer-group 1433
Command Syntax 1433
Parameters 1433
Command Mode 1433
Applicability 1433
Examples 1433
clear bgp ipv6 unicast flap-statistics 1434
Command Syntax 1434
Parameters 1434
Command Mode 1434
Applicability 1434
Examples 1434
clear ipv6 bgp * vrf 1435
Command Syntax 1435
Parameters 1435
Command Mode 1435
Applicability 1435
Examples 1435
clear ipv6 bgp X:X::X:X vrf 1436
Command Syntax 1436
Parameters 1436

Command Mode 1436
Applicability 1436
Examples 1436
clear ip bpg ipv6 unicast table-map 1437
Command Syntax 1437
Parameters 1437
Command Mode 1437
Applicability 1437
Examples 1437
network X:X::X:X 1438
Command Syntax 1438
Parameters 1438
Default 1438
Command Mode 1438
Applicability 1438
Examples 1438
BGP Flowspec Commands 1440
address-family 1441
Command Syntax 1441
Parameters 1441
Default 1441
Command Mode 1441
Applicability 1441
Examples 1441
address-family ipv4 flowspec vrf 1442
Command Syntax 1442
Parameters 1442
Default 1442
Command Mode 1442
Applicability 1442
Examples 1442
class-map type traffic match-all 1443
Command Syntax 1443
Parameters 1443
Default 1443
Command Mode 1443
Applicability 1443
Examples 1443
drop 1444
Command Syntax 1444
Parameters 1444
Default 1444
Command Mode 1444
Applicability 1444
Examples 1444
flowspec 1445
Command Syntax 1445

Parameters 1445
Default 1445
Command Mode 1445
Applicability 1445
Examples 1445
hardware filter ipv4-bgp-flowspec 1446
Command Syntax 1446
Parameters 1446
Default 1446
Command Mode 1446
Applicability 1446
Examples 1446
ipv4 flowspec-disable 1447
Command Syntax 1447
Parameters 1447
Default 1447
Command Mode 1447
Applicability 1447
Examples 1447
local-install 1451
Command Syntax 1451
Parameters 1451
Default 1451
Command Mode 1451
Applicability 1451
Examples 1451
match destination-address 1452
Command Syntax 1452
Parameters 1452
Default 1452
Command Mode 1452
Applicability 1452
Examples 1452
match destination-port 1453
Command Syntax 1453
Parameters 1453
Default 1453
Command Mode 1453
Applicability 1453
Examples 1453
match dscp 1454
Command Syntax 1454
Parameters 1454
Default 1454
Command Mode 1454
Applicability 1454
Examples 1454

© 2026 IP Infusion Inc. Proprietary

match fragment-type 1455
Command Syntax 1455
Parameters 1455
Default 1455
Command Mode 1455
Applicability 1455
Examples 1455
match icmp-code 1456
Command Syntax 1456
Parameters 1456
Default 1456
Command Mode 1456
Applicability 1456
Examples 1456
match icmp-type 1457
Command Syntax 1457
Parameters 1457
Default 1457
Command Mode 1457
Applicability 1457
Examples 1457
match packet-length 1458
Command Syntax 1458
Parameters 1458
Default 1458
Command Mode 1458
Applicability 1458
Examples 1458
match port 1459
Command Syntax 1459
Parameters 1459
Default 1459
Command Mode 1459
Applicability 1459
Examples 1459
match protocol 1460
Command Syntax 1460
Parameters 1460
Default 1460
Command Mode 1460
Applicability 1460
Examples 1460
match source-address 1461
Command Syntax 1461
Parameters 1461
Default 1461
Command Mode 1461

© 2026 IP Infusion Inc. Proprietary

Applicability 1461
Examples 1461
match source-port 1462
Command Syntax 1462
Parameters 1462
Default 1462
Command Mode 1462
Applicability 1462
Examples 1462
match tcp-flag 1463
Command Syntax 1463
Parameters 1463
Default 1463
Command Mode 1463
Applicability 1463
Examples 1463
neighbor flowspec-vaildation-disable 1464
Command Syntax 1464
Parameters 1464
Default 1464
Command Mode 1464
Applicability 1464
Examples 1464
police rate 1465
Command Syntax 1465
Parameters 1465
Default 1465
Command Mode 1465
Applicability 1465
Examples 1465
policy-map 1466
Command Syntax 1466
Parameters 1466
Default 1466
Command Mode 1466
Applicability 1466
Examples 1466
policy-map class type traffic 1467
Command Syntax 1467
Parameters 1467
Default 1467
Command Mode 1467
Applicability 1467
Examples 1467
redirect 1468
Command Syntax 1468
Parameters 1468

© 2026 IP Infusion Inc. Proprietary

Default ... 1468
Command Mode ... 1468
Applicability ... 1468
Examples ... 1468
service-policy type pbr ... 1469
Command Syntax ... 1469
Parameters ... 1469
Default ... 1469
Command Mode ... 1469
Applicability ... 1469
Examples ... 1469
set dscp ... 1470
Command Syntax ... 1470
Parameters ... 1470
Default ... 1470
Command Mode ... 1470
Applicability ... 1470
Examples ... 1470
set ip next-hop ... 1471
Command Syntax ... 1471
Parameters ... 1471
Command Mode ... 1471
Applicability ... 1471
Examples ... 1471
set ipv6 next-hop ... 1472
Command Syntax ... 1472
Parameters ... 1472
Command Mode ... 1472
Applicability ... 1472
Examples ... 1472
show hsl bgp-flowspec-rules fib-id ... 1473
Command Syntax ... 1473
Parameters ... 1473
Default ... 1473
Command Mode ... 1473
Applicability ... 1473
Examples ... 1473
show ip bgp flowspec ... 1475
Command Syntax ... 1475
Parameters ... 1475
Default ... 1475
Command Mode ... 1475
Applicability ... 1475
Examples ... 1475
vrf ... 1481
Command Syntax ... 1481
Parameters ... 1481

© 2026 IP Infusion Inc. Proprietary

Default ... 1481
Command Mode ... 1481
Applicability ... 1481
Examples ... 1481
BGP Virtual Private Network Commands ... 1482
bgp external-route-leak ... 1483
Command Syntax ... 1483
Parameters ... 1483
Default ... 1483
Command Mode ... 1483
Applicability ... 1483
Examples ... 1483
bgp inbound-route-filter ... 1484
Command Syntax ... 1484
Parameter ... 1484
Default ... 1484
Command Mode ... 1484
Applicability ... 1484
Examples ... 1484
clear ip bgp <1-4294967295> vpnv4 ... 1485
Command Syntax ... 1485
Parameters ... 1485
Command Mode ... 1485
Applicability ... 1485
Examples ... 1485
clear ip bgp A.B.C.D vpnv4 ... 1486
Command Syntax ... 1486
Parameters ... 1486
Command Mode ... 1486
Applicability ... 1486
Examples ... 1486
clear bgp <1-4294967295> l2vpn vpls ... 1487
Command Syntax ... 1487
Parameters ... 1487
Command Mode ... 1487
Applicability ... 1487
Example ... 1487
clear bgp A.B.C.D l2vpn vpls ... 1488
Command Syntax ... 1488
Parameters ... 1488
Command Mode ... 1488
Applicability ... 1488
Example ... 1488
clear bgp * l2vpn vpls ... 1489
Command Syntax ... 1489
Parameters ... 1489
Command Mode ... 1489

Applicability 1489
Example 1489
clear ip bgp * vpnv4 1490
Command Syntax 1490
Parameters 1490
Command Mode 1490
Applicability 1490
Examples 1490
debug bgp mpls 1491
Command Syntax 1491
Parameters 1491
Command Mode 1491
Applicability 1491
Examples 1491
export map 1492
Command Syntax 1492
Parameters 1492
Default 1492
Command Mode 1492
Applicability 1492
Examples 1492
import map 1495
Command Syntax 1495
Parameters 1495
Default 1495
Command Mode 1495
Applicability 1495
Examples 1495
ip vrf 1496
Command Syntax 1496
Parameter 1496
Default 1496
Command Mode 1496
Applicability 1496
Command Example 1496
neighbor allow-ebgp-vpn 1497
Command Syntax 1497
Parameters 1497
Default 1497
Command Mode 1497
Applicability 1497
Examples 1497
neighbor as-override 1498
Command Syntax 1498
Parameters 1498
Default 1498
Command Mode 1498

© 2026 IP Infusion Inc. Proprietary

Applicability 1498
Examples 1498
neighbor send-community 1499
Command Syntax 1499
Parameters 1499
Default 1499
Command Mode 1499
Applicability 1499
Examples 1499
neighbor soo 1500
Command Syntax 1500
Parameters 1500
Default 1500
Command Mode 1500
Applicability 1500
Examples 1500
rd (route distinguisher) 1501
Command Syntax 1501
Parameters 1501
Default 1501
Command Mode 1501
Applicability 1501
Examples 1501
route-target 1502
Command Syntax 1502
Parameters 1502
Default 1502
Command Mode 1502
Applicability 1502
Examples 1502
BGP Show Commands 1504
show bgp 1506
Command Syntax 1506
Parameters 1506
Command Mode 1506
Applicability 1506
Examples 1506
show bgp A.B.C.D 1507
Command Syntax 1507
Parameters 1507
Command Mode 1507
Applicability 1507
Examples 1507
show bgp A.B.C.D/M 1509
Command Syntax 1509
Parameters 1509
Command Mode 1509

Applicability 1509
Examples 1509
show bgp client 1511
Command Syntax 1511
Parameters 1511
Command Mode 1511
Applicability 1511
Examples 1511
show bgp community 1512
Command Syntax 1512
Parameters 1512
Command Mode 1512
Applicability 1513
Example 1513
show bgp community-list 1514
Command Syntax 1514
Parameters 1514
Command Mode 1514
Applicability 1514
Example 1514
show bgp dampening dampened-paths 1515
Command Syntax 1515
Parameters 1515
Command Mode 1515
Applicability 1515
Example 1515
show bgp dampening flap-statistics 1517
Command Syntax 1517
Parameters 1517
Command Mode 1517
Applicability 1517
Examples 1517
show bgp dampening parameters 1519
Command Syntax 1519
Parameters 1519
Command Mode 1519
Applicability 1519
Examples 1519
show bgp filter-list 1521
Command Syntax 1521
Parameters 1521
Command Mode 1521
Applicability 1521
Example 1522
show bgp inconsistent-as 1523
Command Syntax 1523
Parameters 1523

Command Mode 1523
Applicability 1523
Examples 1523
show bgp ipv6 1524
Command Syntax 1524
Parameters 1524
Command Mode 1524
Applicability 1524
Example: iBGP and eBGP Routes 1524
Example: IPv6 Prefix Routes 1526
show bgp ipv6 peer-group 1527
Command Syntax 1527
Parameters 1527
Command Mode 1527
Applicability 1527
Example 1527
show bgp l2vpn vpls 1528
Command Syntax 1528
Parameters 1528
Command Mode 1528
Applicability 1528
Example 1528
show bgp label-pool 1531
Command Syntax 1531
Parameters 1531
Command Mode 1531
Applicability 1531
Example 1531
show bgp neighbors 1532
Command Syntax 1532
Parameters 1532
Command Mode 1533
Applicability 1533
Examples 1533
show bgp neighbors advertised-routes 1536
Command Syntax 1536
Parameters 1536
Command Mode 1536
Applicability 1536
Example 1536
show bgp neighbors received prefix-filter 1537
Command Syntax 1537
Parameters 1537
Command Mode 1537
Applicability 1537
Example 1537
show bgp neighbors received-routes 1538

Command Syntax 1538
Parameters 1538
Command Mode 1538
Applicability 1538
Example 1538
show bgp neighbors routes 1540
Command Syntax 1540
Parameters 1540
Command Mode 1540
Applicability 1540
Example 1540
show bgp nexthop-tracking 1542
Command Syntax 1542
Parameters 1542
Command Mode 1542
Applicability 1542
Example 1542
show bgp nexthop-tree-details 1544
Command Syntax 1544
Parameters 1544
Command Mode 1544
Applicability 1544
Example 1544
show bgp paths 1545
Command Syntax 1545
Parameters 1545
Command Mode 1545
Applicability 1545
Example 1545
show bgp prefix-list 1546
Command Syntax 1546
Parameters 1546
Command Mode 1546
Applicability 1546
Example 1546
show bgp quote-regexp 1548
Command Syntax 1548
Parameters 1548
Command Mode 1548
Applicability 1548
Example 1548
show bgp regexp 1549
Command Syntax 1549
Parameters 1549
Command Mode 1549
Applicability 1549
Example 1549

show bgp route-map 1550
Command Syntax 1550
Parameters 1550
Command Mode 1550
Applicability 1550
Example 1550
show bgp statistics 1552
Command Syntax 1552
Parameters 1552
Command Mode 1552
Applicability 1552
Examples 1552
show bgp summary 1554
Command Syntax 1554
Parameters 1554
Command Mode 1554
Applicability 1554
Example 1554
show bgp X:X::X:X 1557
Command Syntax 1557
Parameters 1557
Command Mode 1557
Applicability 1557
Example 1557
show bgp X:X::X:X/M longer prefixes 1558
Command Syntax 1558
Parameters 1558
Command Mode 1558
Applicability 1558
Examples 1558
show debugging bgp 1559
Command Syntax 1559
Parameters 1559
Command Mode 1559
Applicability 1559
Example 1559
show ip bgp 1560
Command Syntax 1560
Parameters 1560
Command Mode 1560
Applicability 1560
Examples 1560
show ip bgp cidr-only 1563
Command Syntax 1563
Parameters 1563
Command Mode 1563
Applicability 1563

Example 1563
show ip bgp community-info 1564
Command Syntax 1564
Parameters 1564
Command Mode 1564
Applicability 1564
Example 1564
show ip bgp peer-group 1565
Command Syntax 1565
Parameters 1565
Command Mode 1565
Applicability 1565
Example 1565
show ip bgp peer-group vrf all 1566
Command Syntax 1566
Parameters 1566
Command Mode 1566
Applicability 1566
Example 1566
show ip bgp rtfilter all 1567
Command Syntax 1567
Parameters 1567
Command Mode 1567
Applicability 1567
Example 1567
show ip bgp rtfilter neighbors 1570
Command Syntax 1570
Parameters 1570
Command Mode 1570
Applicability 1570
Example 1570
show ip bgp scan 1572
Command Syntax 1572
Parameters 1572
Command Mode 1572
Applicability 1572
Example 1572
show ip bgp vpnv4 1573
Command Syntax 1573
Parameters 1573
Command Mode 1574
Applicability 1574
Example 1574
show ip bgp vpnv6 all neighbors 1579
Command Syntax 1579
Parameters 1579
Command Mode 1579

© 2026 IP Infusion Inc. Proprietary

Applicability 1579
Example 1579
show ip bgp vpnv6 rd neighbors 1582
Command Syntax 1582
Parameters 1582
Command Mode 1582
Applicability 1582
Example 1582
show ip extcommunity-list 1584
Command Syntax 1584
Parameters 1584
Command Mode 1584
Applicability 1584
Example 1584
show ip protocols 1585
Command Syntax 1585
Parameters 1585
Command Mode 1585
Applicability 1585
Example 1585
show ip vrf 1586
Command Syntax 1586
Parameters 1586
Default 1586
Command Mode 1586
Applicability 1586
Example 1586
show running-config as-path access-list 1587
Command Syntax 1587
Parameters 1587
Command Mode 1587
Applicability 1587
Example 1587
show running-config bgp 1588
Command Syntax 1588
Parameters 1588
Command Mode 1588
Applicability 1588
Example 1588
show running-config community-list 1589
Command Syntax 1589
Parameters 1589
Command Mode 1589
Applicability 1589
Example 1589
show running-config vrf 1590
Command Syntax 1590

Parameters 1590
Command Mode 1590
Applicability 1590
Example 1590
show running-config ip vrf 1591
Command Syntax 1591
Parameters 1591
Command Mode 1591
Applicability 1591
Example 1591
show running-config mac vrf 1592
Command Syntax 1592
Parameters 1592
Command Mode 1592
Applicability 1592
Example 1592
show running-config vrf word 1593
Command Syntax 1593
Parameters 1593
Command Mode 1593
Applicability 1593
Example 1593
BGP Additional Paths Commands 1594
bgp additional-paths 1595
Command Syntax 1595
Parameters 1595
Default 1595
Command Mode 1595
Applicability 1595
Examples 1595
bgp additional-paths install 1596
Command Syntax 1596
Parameters 1596
Default 1596
Command Mode 1596
Applicability 1596
Examples 1596
neighbor additional-paths 1597
Command Syntax 1597
Parameters 1597
Default 1597
Command Mode 1597
Applicability 1597
Examples 1597
neighbor advertise additional-paths 1598
Command Syntax 1598
Parameters 1598

Default 1598
Command Mode 1598
Applicability 1598
Examples 1598
BGP Graceful Restart Commands 1599
bgp g-shut 1600
Command Syntax 1600
Parameters 1600
Default 1600
Command Mode 1600
Applicability 1600
Examples 1600
bgp g-shut-capable 1601
Command Syntax 1601
Parameters 1601
Default 1601
Command Mode 1601
Applicability 1601
Examples 1601
bgp g-shut-local-preference 1602
Command Syntax 1602
Parameters 1602
Default 1602
Command Mode 1602
Applicability 1602
Examples 1602
bgp graceful-restart 1603
Command Syntax 1603
Parameters 1603
Default 1603
Command Mode 1603
Applicability 1603
Examples 1603
bgp update-delay 1605
Command Syntax 1605
Parameters 1605
Default 1605
Command Mode 1605
Applicability 1605
Example 1605
neighbor capability graceful-restart 1606
Command Syntax 1606
Parameters 1606
Default 1606
Command Mode 1606
Applicability 1606
Examples 1606

© 2026 IP Infusion Inc. Proprietary

neighbor g-shut 1608
Command Syntax 1608
Parameters 1608
Default 1608
Command Mode 1608
Applicability 1608
Examples 1608
neighbor g-shut-timer 1609
Command Syntax 1609
Parameters 1609
Default 1609
Command Mode 1609
Applicability 1609
Examples 1609
neighbor restart-time 1610
Command Syntax 1610
Parameters 1610
Default 1610
Command Mode 1610
Applicability 1610
Example 1610
BGP Labeled Unicast Commands 1611
address-family labeled-unicast 1612
Command Syntax 1612
Parameters 1612
Default 1613
Command Mode 1613
Applicability 1613
Example 1613
allocate-label 1614
Command Syntax 1614
Parameters 1614
Default 1614
Command Mode 1614
Applicability 1614
Example 1614
clear ip bgp * ipv4 labeled-unicast 1615
Command Syntax 1615
Parameters 1615
Command Mode 1615
Applicability 1615
Examples 1615
clear ip bgp A.B.C.D ipv4 labeled-unicast 1616
Command Syntax 1616
Parameters 1616
Command Mode 1616
Applicability 1616

Examples 1616
clear ip bgp peer-group WORD ipv4 labeled-unicast 1617
Command Syntax 1617
Parameters 1617
Command Mode 1617
Applicability 1617
Examples 1617
Regular Expressions 1618

Bidirectional Forwarding Detection Configuration 1619

BFD Authentication 1621
Overview 1621
Enabling BFD Authentication for Multihop for IPv4, Using Key-ID 1621
Topology 1621
Configuration 1622
R1 Configuration 1622
R2 Configuration 1622
R3 Configuration 1623
Validation 1623
Base BFD Configuration 1626
Overview 1626
Topology 1626
BFD Echo Function 1627
BFD Slow Timer 1629
BFD Multihop Peer Timer 1630
Validation: 1631
BFD Single-hop Session Timer 1633
Validation: 1633
BFD Echo Interval 1635
BFD Protocol Configurations 1637
OSPF—BFD Single-Hop Session 1637
Topology 1637
Validation 1638
OSPF—BFD Multi-Hop Session 1644
Topology 1644
Validation 1645
BFD Configuration in IS-IS 1647
Topology 1647
Validation 1648
BFD Configuration in BGP 1649
Topology 1650
Validation 1652
BFD Static Route Configuration 1657
Topology 1657
R1 1657
R2 1657
R3 1658

R4 1658
Validation 1658
R3 1659
BFD with VRF Configuration 1661
Topology 1661
RTR1 Configuration 1661
Switch Configuration 1662
RTR2 Configuration 1663
BFD with VRF Validation 1664
RTR1 1664
RTR3 1665
BFD Over Static Routing IPv4 and IPv6 1667
RTR1 Configuration 1667
RTR2 Configuration 1667
BFD Over Static Routing Validation 1668
BFD Support on LAG Interface 1671
Overview 1671
Single Hop BFD Option Over LAG 1671
Feature Characteristics 1671
Benefits 1671
Prerequisites 1672
Configuration 1672
Validation 1673
CLI Commands 1674
bfd session software 1674
Glossary 1675

Bidirectional Forwarding Detection Command Reference 1676

Bidirectional Forwarding Commands 1678
accept-lifetime 1679
Command Syntax 1679
Parameters 1679
Default 1679
Command Mode 1679
Applicability 1680
Examples 1680
bfd auth type 1681
Command Syntax 1681
Parameters 1681
Command Mode 1681
Applicability 1681
Example 1682
bfd 1683
Command Syntax 1683
Parameters 1683
Command Mode 1683
Applicability 1683

Example 1683
bfd echo 1684
Command Syntax 1684
Parameters 1684
Default 1684
Command Mode 1684
Applicability 1684
Example 1684
bfd echo interval 1685
Command Syntax 1685
Parameter 1685
Default 1685
Command Mode 1685
Applicability 1685
Example 1685
bfd echo ipv4 source 1686
Command Syntax 1686
Parameter 1686
Command Mode 1686
Applicability 1686
Example 1686
bfd interval 1687
Command Syntax 1687
Parameters 1687
Defaults 1687
Command Mode 1687
Applicability 1687
Example 1687
bfd multihop-peer 1688
Command Syntax 1688
Parameters 1688
Command Mode 1688
Applicability 1689
Example 1689
bfd multihop-peer A.B.C.D interval 1690
Command Syntax 1690
Parameters 1690
Command Mode 1690
Default 1690
Applicability 1690
Example 1690
bfd multihop-peer interval 1691
Command Syntax 1691
Parameters 1691
Command Mode 1691
Default 1691
Applicability 1691

Example ... 1691
bfd multihop-peer X:X::X:X interval ... 1692
Command Syntax ... 1692
Parameters ... 1692
Command Mode ... 1692
Default ... 1692
Applicability ... 1692
Example ... 1692
bfd notification ... 1693
Command Syntax ... 1693
Parameters ... 1693
Command Mode ... 1693
Applicability ... 1693
Examples ... 1693
bfd session ... 1694
Command Syntax ... 1694
Parameters ... 1694
Command Mode ... 1694
Applicability ... 1694
Examples ... 1694
bfd slow-timer ... 1696
Command Syntax ... 1696
Parameter ... 1696
Command Mode ... 1696
Default ... 1696
Applicability ... 1696
Example ... 1696
debug bfd ... 1697
Command Syntax ... 1697
Parameters ... 1697
Command Mode ... 1697
Applicability ... 1697
Examples ... 1697
hardware-profile seamless-bfd ... 1698
Command Syntax ... 1698
Parameter ... 1698
Default ... 1698
Command Mode ... 1698
Applicability ... 1698
Example ... 1698
key-id ... 1699
Command Syntax ... 1699
Parameters ... 1699
Default ... 1699
Command Mode ... 1699
Applicability ... 1699
Examples ... 1699

key chain 1700
Command Syntax 1700
Parameters 1700
Command Mode 1700
Applicability 1700
Examples 1700
key-string 1701
Command Syntax 1701
Parameters 1701
Default 1701
Command Mode 1701
Applicability 1701
Examples 1701
key-string encrypted 1702
Command Syntax 1702
Parameters 1702
Default 1702
Command Mode 1702
Applicability 1702
Example 1702
send-lifetime 1703
Command Syntax 1703
Parameters 1703
Command Mode 1703
Applicability 1703
Examples 1704
show bfd 1705
Command Syntax 1705
Parameters 1705
Command Mode 1705
Applicability 1705
Example 1705
show bfd interface 1706
Command Syntax 1706
Parameters 1706
Command Mode 1706
Applicability 1706
Example 1706
show bfd session 1708
Command Syntax 1708
Parameters 1708
Command Mode 1708
Applicability 1708
Example 1708
show bfd session A.B.C.D 1713
Command Syntax 1713
Parameters 1713

© 2026 IP Infusion Inc. Proprietary

Command Mode 1713
Applicability 1713
Example 1713
show bfd session ipv6 1716
Command Syntax 1716
Parameters 1716
Command Mode 1716
Applicability 1716
Example 1716
show debugging bfd 1719
Command Syntax 1719
Parameters 1719
Command Mode 1719
Applicability 1719
Example 1719
snmp restart bfd 1720
Command Syntax 1720
Parameters 1720
Command Mode 1720
Applicability 1720
Example 1720
Protocol Commands for BFD 1721
area virtual-link 1722
Command Syntax 1722
Parameters 1722
Default 1722
Command Mode 1722
Applicability 1722
Example 1722
bfd all-interfaces 1723
Command Syntax 1723
Parameters 1723
Default 1723
Command Mode 1723
Applicability 1723
Example 1723
debug bgp bfd 1724
Command Syntax 1724
Parameters 1724
Command Mode 1724
Applicability 1724
Example 1724
debug isis bfd 1725
Command Syntax 1725
Parameters 1725
Command Mode 1725
Applicability 1725

© 2026 IP Infusion Inc. Proprietary

Example 1725
debug ospf bfd 1726
Command Syntax 1726
Parameters 1726
Command Mode 1726
Applicability 1726
Example 1726
ip ospf bfd 1727
Command Syntax 1727
Parameter 1727
Default 1727
Command Mode 1727
Applicability 1727
Example 1727
isis bfd 1728
Command Syntax 1728
Parameter 1728
Default 1728
Command Mode 1728
Applicability 1728
Example 1728
BFD Static Route Commands 1729
ip bfd static all-interfaces 1730
Command Syntax 1730
Parameters 1730
Command Mode 1730
Applicability 1730
Example 1730
ip static fall-over-bfd 1731
Command Syntax 1731
Parameters 1731
Command Mode 1731
Applicability 1731
Example 1731
ip static bfd 1732
Command Syntax 1732
Parameters 1732
Command Mode 1732
Default 1732
Applicability 1732
Example 1732
ipv6 bfd static all-interfaces 1733
Command Syntax 1733
Parameters 1733
Command Mode 1733
Applicability 1733
Example 1733

ipv6 static fall-over-bfd 1734
Command Syntax 1734
Parameters 1734
Command Mode 1734
Applicability 1734
Examples 1734
ipv6 static bfd 1735
Command Syntax 1735
Parameters 1735
Command Mode 1735
Applicability 1735
Example 1735

Intermediate System to Intermediate System Configuration 1736

IS-IS IPv4 1738
Enable IS-IS on an Interface Configuration 1738
Topology 1738
R1 Configuration 1738
R2 Configuration 1738
Validation 1739
Set Priority 1741
Topology 1741
R1 Configuration 1741
R2 Configuration 1741
Validation 1742
Dynamic Hostname 1743
Topology 1743
R1 Configuration 1743
R2 Configuration 1744
Validation 1744
Redistribute Routes into IS-IS 1745
Topology 1745
R1 Configuration 1745
R2 Configuration 1746
R3 Configuration 1747
Validation 1747
Metric 1750
Topology 1751
R1 Configuration 1751
R2 Configuration 1752
R3 Configuration 1752
R4 Configuration 1753
Validation 1754
L1-L2 Area Routing with a Single Instance 1758
Topology 1758
R1 Configuration 1758
R2 Configuration 1759

R3 Configuration 1759
R4 Configuration 1760
Validation 1760
L1-L2 Area Routing with Multiple Instances 1764
Topology 1764
R1 Configuration 1764
R2 Configuration 1765
R3 Configuration 1766
R4 Configuration 1766
Validation 1767
Route Leaking 1769
Topology 1770
R1 Configuration 1770
R2 Configuration 1770
R3 Configuration 1771
Validation 1772
Route Summarization 1774
Topology 1775
R1 Configuration 1775
R2 Configuration 1775
R3 Configuration 1776
Validation 1776
IS-IS Distance 1779
Topology 1779
R1 Configuration 1779
R2 Configuration 1780
R3 Configuration 1781
Validation 1781
Overload Bit 1785
Overview 1785
On-Startup Wait-for-BGP 1785
R3 1795
R4 1795
R3 1795
Suppress External and Inter-level Reachability 1796
Passive Interface 1805
Usage 1805
Topology 1806
R1 Configuration 1806
R2 Configuration 1807
R3 Configuration 1807
Validation 1808
IS-IS IPv4 Loop-Free Alternate Fast Reroute 1810
Overview 1810
Basic Configuration 1811
Backup Path Based on Route-Map Prefixes 1818
LFA Tie-Breaker 1819

© 2026 IP Infusion Inc. Proprietary

LFA Termination ... 1821
LFA For ECMP Paths ... 1824
Validation ... 1831
IS-IS IPv6 Configuration ... 1833
Enable IS-ISv6 on an Interface ... 1833
Configuration ... 1833
Validation ... 1834
Set Priority ... 1835
Configuration ... 1835
Validation ... 1836
Dynamic hostname ... 1837
Configuration ... 1838
Validation ... 1839
Redistribute Routes into IS-IS ... 1840
Topology ... 1840
Configuration ... 1840
Validation ... 1842
Interface Metric ... 1842
Topology ... 1843
Configuration ... 1843
Validation ... 1846
Route Summarization ... 1847
Topology ... 1847
Configuration ... 1847
Validation ... 1849
Passive Interface ... 1852
Topology ... 1853
Configuration ... 1853
Validation ... 1854
Enable BFD over IS-ISv6 ... 1857
Topology ... 1857
Configuration ... 1857
Validation ... 1858
Originate Default Route to ISISv6 Neighbors ... 1859
Topology ... 1859
Configuration ... 1860
Validation ... 1861
IS-IS Distance ... 1862
Topology ... 1862
Configuration ... 1862
Validation ... 1864
IS-IS IPv4 TE-Metric Extension ... 1865
Overview ... 1865
Topology ... 1865
Configuration ... 1865
R1 ... 1865
R2 ... 1866

© 2026 IP Infusion Inc. Proprietary

Validation ... 1866
IS-IS Graceful Restart Configuration ... 1870
Overview ... 1870
Topology ... 1870
Configuration ... 1870
Validation ... 1871
IS-IS-TE IPv4 ... 1873
Enable MPLS-TE in Level-1 and Level-2 on L1-L2 IS ... 1873
Toplogy ... 1873
Configuration ... 1873
Validation ... 1875
Maximum Link Bandwidth and Reservable Bandwidth ... 1879
Topology ... 1879
Configuration ... 1879
Validation ... 1882
Administrative Group Constraints ... 1887
Configuration ... 1887
Validation ... 1890
ISIS Multi Topology ... 1895
Overview ... 1895
Feature Characteristics ... 1895
Benefits ... 1895
Prerequisites ... 1895
Configuration ... 1898
Topology ... 1898
Validation for Multi Topology ... 1899
R1 ... 1899
R2 ... 1904
R3 ... 1909
R4 ... 1914
R5 ... 1919
Running Configuration ... 1923
CLI Commands ... 1923
multi topology ... 1923
Command Syntax ... 1923
Parameters ... 1923
Default ... 1923
Command Mode ... 1924
Applicability ... 1924
Example ... 1924
Glossary ... 1924

Intermediate System to Intermediate System Command Reference ... 1925

IS-IS Commands ... 1929
accept-lifetime ... 1932
Command Syntax ... 1932
Parameters ... 1932

Default 1932
Command Mode 1932
Applicability 1933
Examples 1933
address-family ipv6 1934
Command Syntax 1934
Parameters 1934
Default 1934
Command Mode 1934
Example 1934
adjacency-check 1935
Command Syntax 1935
Parameters 1935
Default 1935
Command Mode 1935
Applicability 1935
Examples 1935
area-password 1936
Command Syntax 1936
Parameters 1936
Default 1936
Command Mode 1936
Applicability 1936
Examples 1936
authentication key-chain 1937
Command Syntax 1937
Parameters 1937
Default 1937
Command Mode 1937
Applicability 1937
Example 1937
authentication mode 1938
Command Syntax 1938
Parameters 1938
Default 1938
Command Mode 1938
Applicability 1938
Examples 1938
authentication send-only 1939
Command Syntax 1939
Parameters 1939
Default 1939
Command Mode 1939
Applicability 1939
Examples 1939
bfd all-interfaces 1940
Command Syntax 1940

Parameters 1940
Default 1940
Command Mode 1940
Applicability 1940
Examples 1940
capability cspf 1941
Command Syntax 1941
Parameters 1941
Default 1941
Command Mode 1941
Example 1941
clear clns is-neighbors 1942
Command Syntax 1942
Parameters 1942
Command Mode 1942
Applicability 1942
Example 1942
clear clns neighbors 1943
Command Syntax 1943
Parameters 1943
Command Mode 1943
Applicability 1943
Example 1943
clear ip isis route 1944
Command Syntax 1944
Parameters 1944
Command Mode 1944
Applicability 1944
Example 1944
clear isis adjacency 1945
Command Syntax 1945
Parameters 1945
Command Mode 1945
Applicability 1945
Example 1945
clear isis counter 1946
Command Syntax 1946
Parameters 1946
Command Mode 1946
Applicability 1946
Example 1946
clear isis interface counter 1947
Command Syntax 1947
Parameters 1947
Command Mode 1947
Applicability 1947
Example 1947

© 2026 IP Infusion Inc. Proprietary

clear isis process 1948
Command Syntax 1948
Parameters 1948
Command Mode 1948
Applicability 1948
Example 1948
debug isis 1949
Command syntax 1949
Parameters 1949
Command Mode 1950
Applicability 1950
Examples 1950
default-information originate 1951
Command Syntax 1951
Parameters 1951
Default 1951
Command Mode 1951
Applicability 1951
Examples 1951
distance (IPv4) 1952
Command Syntax 1952
Parameters 1952
Default 1952
Command Mode 1952
Applicability 1952
Examples 1952
distance (IPv6) 1954
Command Syntax 1954
Parameters 1954
Default 1954
Command Mode 1954
Applicability 1954
Examples 1954
domain password 1956
Command Syntax 1956
Parameters 1956
Default 1956
Command Mode 1956
Applicability 1956
Examples 1956
dynamic-hostname 1958
Command Syntax 1958
Parameters 1958
Default 1958
Command Mode 1958
Example 1958
fast-reroute per-prefix 1959

Command Syntax 1959
Parameters 1959
Command Mode 1959
Applicability 1959
Examples 1959
fast-reroute per-prefix remote-lfa 1960
Command Syntax 1960
Parameters 1960
Default 1960
Command Mode 1960
Applicability 1960
Examples 1960
fast-reroute terminate-hold-on interval 1961
Command Syntax 1961
Parameters 1961
Default 1961
Command Mode 1961
Applicability 1961
Examples 1961
fast-reroute tie-break 1962
Command Syntax 1962
Parameters 1962
Defaults 1963
Command Mode 1963
Applicability 1963
Examples 1963
ignore-lsp-errors 1964
Command Syntax 1964
Parameters 1964
Default 1964
Command Mode 1964
Applicability 1964
Example 1964
ip route high-priority tag 1965
Command Syntax 1965
Parameters 1965
Default 1965
Command Mode 1965
Applicability 1965
Example 1965
ip router isis 1966
Command Syntax 1966
Parameters 1966
Default 1966
Command Mode 1966
Applicability 1966
Example 1966

© 2026 IP Infusion Inc. Proprietary

ipv6 router isis 1967
Command Syntax 1967
Parameters 1967
Default 1967
Command Mode 1967
Example 1967
is-type 1968
Command Syntax 1968
Parameters 1968
Default 1968
Command Mode 1968
Applicability 1968
Examples 1968
isis authentication key-chain 1969
Command Syntax 1969
Parameters 1969
Default 1969
Command Mode 1969
Applicability 1969
Example 1969
isis authentication mode md5 1970
Command Syntax 1970
Parameters 1970
Default 1970
Command Mode 1970
Applicability 1970
Example 1970
isis authentication send-only 1971
Command Syntax 1971
Parameters 1971
Default 1971
Command Mode 1971
Applicability 1971
Example 1971
isis bfd 1972
Command Syntax 1972
Parameters 1972
Default 1972
Command Mode 1972
Applicability 1972
Example 1972
isis circuit-type 1973
Command Syntax 1973
Parameters 1973
Default 1973
Command Mode 1973
Applicability 1973

© 2026 IP Infusion Inc. Proprietary

Example 1973
isis csnp-interval 1974
Command Syntax 1974
Parameters 1974
Default 1974
Command Mode 1974
Applicability 1974
Example 1974
isis fast-reroute per-prefix candidate disable 1975
Command Syntax 1975
Parameters 1975
Command Mode 1975
Applicability 1975
Example 1975
isis hello-interval 1976
Command Syntax 1976
Parameters 1976
Default 1976
Command Mode 1976
Applicability 1976
Examples 1977
isis hello-multiplier 1978
Command Syntax 1978
Parameters 1978
Default 1978
Command Mode 1978
Applicability 1978
Example 1978
isis hello padding 1979
Command Syntax 1979
Parameters 1979
Default 1979
Command Mode 1979
Applicability 1979
Examples 1979
isis lsp-interval 1980
Command Syntax 1980
Parameters 1980
Default 1980
Command Mode 1980
Applicability 1980
Examples 1980
isis mesh-group 1981
Command Syntax 1981
Parameters 1981
Default 1981
Command Mode 1981

© 2026 IP Infusion Inc. Proprietary

Examples 1981
isis metric 1982
Command Syntax 1982
Parameters 1982
Default 1982
Command Mode 1982
Applicability 1982
Example 1982
isis network 1983
Command Syntax 1983
Parameters 1983
Default 1983
Command Mode 1983
Applicability 1983
Example 1983
isis password 1984
Command Syntax 1984
Parameters 1984
Default 1984
Command Mode 1984
Applicability 1984
Example 1984
isis priority 1985
Command Syntax 1985
Parameters 1985
Default 1985
Command Mode 1985
Applicability 1985
Example 1985
isis retransmit-interval 1986
Command Syntax 1986
Parameters 1986
Default 1986
Command Mode 1986
Applicability 1986
Examples 1986
isis tag 1987
Command Syntax 1987
Parameters 1987
Command Mode 1987
Examples 1987
isis te-metric 1988
Command Syntax 1988
Parameters 1988
Default 1988
Command Mode 1988
Applicability 1988

Examples 1988
isis wait-timer 1989
Command Syntax 1989
Parameters 1989
Default 1989
Command Mode 1989
Applicability 1989
Examples 1989
isis wide-metric 1990
Command Syntax 1990
Parameters 1990
Default 1990
Command Mode 1990
Applicability 1990
Examples 1990
ispf 1991
Command Syntax 1991
Parameters 1991
Default 1991
Command Mode 1991
Applicability 1991
Examples 1991
key chain 1992
Command Syntax 1992
Parameters 1992
Default 1992
Command Mode 1992
Applicability 1992
Examples 1992
key-id 1993
Command Syntax 1993
Parameters 1993
Default 1993
Command Mode 1993
Applicability 1993
Examples 1993
key-string 1994
Command Syntax 1994
Parameters 1994
Default 1994
Command Mode 1994
Applicability 1994
Examples 1994
key-string encrypted 1995
Command Syntax 1995
Parameters 1995
Default 1995

© 2026 IP Infusion Inc. Proprietary

Command Mode 1995
Applicability 1995
Example 1995
lsp-gen-interval 1996
Command Syntax 1996
Parameters 1996
Default 1996
Command Mode 1996
Applicability 1996
Example 1996
lsp-mtu 1997
Command Syntax 1997
Parameters 1997
Default 1997
Command Mode 1997
Applicability 1997
Example 1997
lsp-refresh-interval 1998
Command Syntax 1998
Parameters 1998
Default 1998
Command Mode 1998
Applicability 1998
Examples 1998
max-area-address 1999
Command Syntax 1999
Parameters 1999
Default 1999
Command Mode 1999
Applicability 1999
Examples 1999
max-lsp-lifetime 2000
Command Syntax 2000
Parameters 2000
Default 2000
Command Mode 2000
Applicability 2000
Examples 2000
metric-style 2001
Command Syntax 2001
Parameters 2001
Default 2001
Command Mode 2002
Examples 2002
mpls traffic-eng 2003
Command Syntax 2003
Parameters 2003

© 2026 IP Infusion Inc. Proprietary

Default 2003
Command Mode 2003
Examples 2003
mpls traffic-eng router-id 2004
Command Syntax 2004
Parameters 2004
Default 2004
Command Mode 2004
Examples 2004
net 2005
Command Syntax 2005
Parameters 2005
Default 2005
Command Mode 2005
Applicability 2005
Example 2005
passive-interface 2006
Command Syntax 2006
Parameters 2006
Command Mode 2006
Examples 2006
prc-interval-exp 2007
Command Syntax 2007
Parameters 2007
Default 2007
Command Mode 2007
Applicability 2007
Examples 2007
redistribute 2008
Command Syntax 2008
Parameters 2008
Default 2009
Command Mode 2009
Applicability 2009
Examples 2009
redistribute isis 2010
Command Syntax 2010
Parameters 2010
Default 2010
Command Mode 2010
Applicability 2010
Examples 2010
redistribute isis WORD 2012
Command Syntax 2012
Parameters 2012
Default 2012
Command Mode 2012

© 2026 IP Infusion Inc. Proprietary

Applicability 2013
Examples 2013
router isis 2014
Command Syntax 2014
Parameters 2014
Default 2014
Command Mode 2014
Applicability 2014
Example 2014
send-lifetime 2015
Command Syntax 2015
Parameters 2015
Default 2015
Command Mode 2015
Applicability 2016
Examples 2016
set-overload-bit 2017
Command Syntax 2017
Parameters 2017
Default 2017
Command Mode 2017
Applicability 2017
Example 2017
snmp context 2019
Command Syntax 2019
Parameters 2019
Default 2019
Command Mode 2019
Applicability 2019
Examples 2019
spf-interval-exp 2020
Command Syntax 2020
Parameters 2020
Default 2020
Command Mode 2020
Applicability 2020
Examples 2020
summary-address 2021
Command Syntax 2021
Parameters 2021
Default 2021
Command Mode 2021
Applicability 2021
Examples 2021
summary-prefix 2022
Command Syntax 2022
Parameters 2022

Default 2022
Command Mode 2022
Applicability 2022
Examples 2022
IS-IS Graceful Restart Commands 2023
capability restart graceful 2024
Command Syntax 2024
Parameters 2024
Default 2024
Command Mode 2024
Applicability 2024
Example 2024
isis restart grace-period 2025
Command Syntax 2025
Parameters 2025
Default 2025
Command Mode 2025
Applicability 2025
Examples 2025
isis restart-hello-interval 2026
Command Syntax 2026
Parameters 2026
Default 2026
Command Mode 2026
Applicability 2026
Example 2026
isis restart helper 2027
Command Syntax 2027
Parameters 2027
Default 2027
Command Mode 2027
Applicability 2027
Example 2027
isis restart suppress-adjacency 2028
Command Syntax 2028
Parameters 2028
Default 2028
Command Mode 2028
Applicability 2028
Example 2028
restart isis graceful 2029
Command Syntax 2029
Parameters 2029
Default 2029
Command Mode 2029
Applicability 2029
Example 2029

© 2026 IP Infusion Inc. Proprietary

restart-timer 2030
Command Syntax 2030
Parameters 2030
Default 2030
Command Mode 2030
Applicability 2030
Example 2030
ISIS Show Commands 2031
show clns is-neighbors 2032
Command Syntax 2032
Parameters 2032
Command Mode 2032
Applicability 2032
Example 2032
show clns neighbors 2034
Command Syntax 2034
Parameters 2034
Command Mode 2034
Applicability 2034
Example 2034
show cspf lsp 2035
Command Syntax 2035
Parameters 2035
Command Mode 2036
Applicability 2036
Example 2036
show cspf rsvp forwarding-timer 2038
Command Syntax 2038
Parameters 2038
Command Mode 2038
Applicability 2038
Example 2038
show debugging isis 2039
Command Syntax 2039
Parameters 2039
Command Mode 2039
Applicability 2039
Example 2039
show ip isis igp-shortcut-lsp 2040
Command Syntax 2040
Parameters 2040
Command Mode 2040
Applicability 2040
Example 2040
show ip isis lfa-config 2041
Command Syntax 2041
Parameters 2041

Command Mode 2041
Default 2041
Applicability 2041
Example 2041
show ip isis route 2042
Command Syntax 2042
Parameters 2042
Command Mode 2042
Applicability 2042
Example 2042
show ip isis route fast-reroute 2044
Command Syntax 2044
Parameters 2044
Command Mode 2044
Applicability 2044
Example 2044
show ip isis route igp-shortcut 2046
Command Syntax 2046
Parameters 2046
Command Mode 2046
Applicability 2046
Example 2046
show ip protocols 2048
Command Syntax 2048
Parameters 2048
Command Mode 2048
Applicability 2048
Example 2048
show ip route fast-reroute 2050
Command Syntax 2050
Parameters 2050
Command Mode 2050
Applicability 2050
Example 2050
show ipv6 isis route 2051
Command Syntax 2051
Parameters 2051
Command Mode 2051
Applicability 2051
Example 2051
show ipv6 isis topology 2053
Command Syntax 2053
Parameters 2053
Command Mode 2053
Applicability 2053
Example 2053
show isis counter 2055

© 2026 IP Infusion Inc. Proprietary

Command Syntax 2055
Parameters 2055
Command Mode 2055
Applicability 2055
Example 2055
show isis database 2056
Command Syntax 2056
Parameters 2056
Command Mode 2056
Applicability 2056
Example 2056
show isis interface 2059
Command Syntax 2059
Parameters 2059
Command Mode 2059
Applicability 2059
Example 2059
show isis microloop-avoidance 2062
Command Syntax 2062
Parameters 2062
Command Mode 2062
Applicability 2062
Example 2062
show isis spf-logs 2064
Command Syntax 2064
Parameters 2064
Command Mode 2064
Applicability 2064
Example 2064
show isis topology 2065
Command Syntax 2065
Parameters 2065
Command Mode 2065
Applicability 2065
Example 2065
show running-config interface isis 2067
Command Syntax 2067
Parameters 2067
Command Mode 2067
Applicability 2067
Example 2067
show running-config router isis 2068
Command Syntax 2068
Parameters 2068
Command Mode 2068
Applicability 2068
Example 2068

© 2026 IP Infusion Inc. Proprietary

IS-IS Microloop-Avoidance Commands 2069
microloop-avoidance 2070
Command Syntax 2070
Parameters 2070
Default 2070
Command Mode 2070
Applicability 2070
Examples 2071
microloop-avoidance flex-algo 2072
Command Syntax 2072
Parameters 2072
Default 2072
Command Mode 2072
Applicability 2072
Examples 2072
microloop-avoidance hold-timer 2073
Command Syntax 2073
Parameters 2073
Default 2073
Command Mode 2073
Applicability 2073
Examples 2073
microloop-avoidance max-fib 2074
Command Syntax 2074
Parameters 2074
Default 2074
Command Mode 2074
Applicability 2074
Examples 2074
microloop-avoidance segment-routing 2075
Command Syntax 2075
Parameters 2075
Default 2075
Command Mode 2075
Applicability 2075
Examples 2075
show isis microloop-avoidance 2076
Command Syntax 2076
Parameters 2076
Command Mode 2076
Applicability 2076
Example 2076

Layer 3 Sub-interface Configuration 2078

Layer 3 Sub-interface Configuration 2079
Topology 2079
Creating a Sub-interface 2079

Creating a Sub-interface with Encapsulation 2080
Displaying Sub-interfaces 2081

Layer 3 Sub-interface Commands 2084
Layer 3 Sub-interface Commands 2085
encapsulation 2086
Command Syntax 2086
Parameters 2086
Command Mode 2086
Applicability 2086
Example 2086
interface IFNAME.SUBINTERFACE_ID 2087
Command Syntax 2087
Parameters 2087
Command Mode 2087
Applicability 2087
Example 2087
show interface IFNAME.SUBINTERFACE_ID 2088
Command Syntax 2088
Parameters 2088
Command Mode 2088
Applicability 2088
Example 2088

Neighbor Discovery Configuration 2090
Neighbor Discovery Configuration 2091
ARP/Neighbor Discovery Operation 2091
Neighbor Entry States 2091
Configuring ARP for IPv4 2092
Configuring Neighbor Discovery for IPv6 2093

Neighbor Discovery Command Reference 2095
Neighbor Discovery Command Reference 2096
arp-ageing-timeout 2097
Command Syntax 2097
Parameter 2097
Default value 2097
Command Mode 2097
Applicability 2097
Example 2097
arp-reachable-time 2098
Command Syntax 2098
Parameters 2098
Default value 2098
Command Mode 2098
Applicability 2098
Example 2098
clear arp 2099

© 2026 IP Infusion Inc. Proprietary

Command Syntax 2099
Parameters 2099
Command Mode 2099
Applicability 2099
clear ipv6 neighbors 2100
Command Syntax 2100
Parameters 2100
Command Mode 2100
Applicability 2100
Example 2100
debug ip arp 2101
Command Syntax 2101
Parameters 2101
Command Mode 2101
Applicability 2101
Example 2101
debug ipv6 nd 2102
Command Syntax 2102
Parameters 2102
Command Mode 2102
Applicability 2102
Example 2102
ip arp 2103
Command Syntax 2103
Parameters 2103
Default value 2103
Command Mode 2103
Applicability 2103
Example 2103
ip arp vrf 2104
Command Syntax 2104
Parameter 2104
Default value 2104
Command Mode 2104
Applicability 2104
Example 2104
ip proxy-arp 2105
Command Syntax 2105
Parameter 2105
Default value 2105
Command Mode 2105
Applicability 2105
Examples 2105
ipv6 nd current-hoplimit 2106
Command Syntax 2106
Parameter 2106
Default 2106

Command Mode 2106
Applicability 2106
Examples 2106
ipv6 nd link-mtu 2107
Command Syntax 2107
Parameters 2107
Default 2107
Command Mode 2107
Applicability 2107
Example 2107
ipv6 nd managed-config-flag 2108
Command Syntax 2108
Parameters 2108
Default 2108
Command Mode 2108
Applicability 2108
Examples 2108
ipv6 nd other-config-flag 2109
Command Syntax 2109
Parameters 2109
Default 2109
Command Mode 2109
Applicability 2109
Examples 2109
ipv6 nd prefix 2110
Command Syntax 2110
Parameters 2110
Default 2110
Command Mode 2110
Applicability 2110
Defaults 2110
Examples 2110
ipv6 nd ra-interval 2112
Command Syntax 2112
Parameter 2112
Default 2112
Command Mode 2112
Applicability 2112
Examples 2112
ipv6 nd ra-lifetime 2113
Command Syntax 2113
Parameter 2113
Default 2113
Command Mode 2113
Applicability 2113
Examples 2113
ipv6 nd reachable-time 2114

Command Syntax 2114
Parameter 2114
Default 2114
Command Mode 2114
Applicability 2114
Examples 2114
ipv6 nd retransmission-time 2115
Command Syntax 2115
Parameter 2115
Default 2115
Command Mode 2115
Applicability 2115
Examples 2115
ipv6 nd suppress-ra 2116
Command Syntax 2116
Parameters 2116
Default 2116
Command Mode 2116
Applicability 2116
Example 2116
ipv6 neighbor 2117
Command Syntax 2117
Parameter 2117
Default value 2117
Command Mode 2117
Applicability 2117
Examples 2117
nd-ageing-timeout 2118
Command Syntax 2118
Parameters 2118
Default value 2118
Command Mode 2118
Applicability 2118
Example 2118
nd-reachable-time 2119
Command Syntax 2119
Parameter 2119
Default value 2119
Command Mode 2119
Applicability 2119
Examples 2119
no debug all 2120
Command Syntax 2120
Parameters 2120
Command Mode 2120
Applicability 2120
Example 2120

© 2026 IP Infusion Inc. Proprietary

show arp 2121
Command Syntax 2121
Parameters 2121
Command Mode 2121
Applicability 2121
Example 2121
show debugging ip arp 2123
Command Syntax 2123
Parameters 2123
Command Mode 2123
Applicability 2123
Example 2123
show debugging ipv6 nd 2124
Command Syntax 2124
Parameters 2124
Command Mode 2124
Applicability 2124
Example 2124
show ipv6 neighbors 2125
Command Syntax 2125
Parameters 2125
Command Mode 2125
Applicability 2125
Example 2125

Open Shortest Path First Configuration 2128

OSPFv2 2130
Enable OSPF on an Interface 2130
Topology 2130
Validation 2131
Set Priority 2134
Topology 2134
Validation 2135
Area Border Router 2140
Topology 2140
Configuration 2140
Validation 2141
Redistribute Routes into OSPF 2144
Topology 2144
Validation 2144
Cost 2146
Topology 2146
Configuration 2146
Validation 2148
Path MTU for IPv4 and IPv6 2152
Topology 2152
Validation 2153

© 2026 IP Infusion Inc. Proprietary

Virtual Links 2155
Topology 2155
Validation 2156
OSPF Authentication 2160
Topology 2160
Validation 2162
Multiple OSPF Instances 2165
Topology 2165
Enable Multiple OSPF Instances on a Router 2165
Validation 2166
Redistribute among Multiple Instances 2171
Redistribute with the Metric Option 2172
Redistribute with the Type Option 2173
Redistribute with the Route-Map Option 2174
Multiple OSPF Instances on Same Subnet 2177
Topology 2177
Configuration 2177
Validation 2178
Multi-Area Adjacency Configuration 2180
Topology 2180
Configuration 2180
Validation 2181
LSA Throttling 2184
How OSPF LSA Throttling Works 2184
Topology 2184
Validation 2185
Configure OSPF LSA Arrival Timers 2190
Loop-Free Alternate Fast Reroute 2196
Overview 2196
Topology 2196
Validation 2199
LFA Tie-Breaker 2202
LFA Termination 2204
Loop-Free Alternate (LFA) ECMP PATH 2206
Overview 2206
Topology 2206
Configuring OSPF LFA ECMP 2207
Validation 2210
Configuring OSPF LFA ECMP with Load Balancing 2213
SNMP Support for Multiple Instance of OSPF Configuration 2214
Overview 2214
Topology 2214
Enable SNMP and create SNMP Context & Group for OSPF Instances on a Router with SNMPv2 2215
Enable SNMP and create SNMP Context & Group for OSPF Instances on a Router with SNMPv3 Configuration 2221
Enable SNMP and create SNMP Context & Group for Multiple OSPF Instances on a 2227

Router with SNMPv2/v3
R2 2232
R3 2233
SNMP WALK Command 2235
OSPF Stub Router Advertisement 2236
Topology 2236
Configuration 2236
Validation 2238
Multi-Area Redundant Adjacency Configuration 2243
Overview 2243
Feature Characteristics 2243
Benefits 2243
Prerequisites 2244
Topology 2244
Configuration 2245
Sample show running-config Output 2246
R1 2246
R2 2247
R3 2247
R4 2247
R5 2248
Validation 2248
R1 2248
R2 2250
R3 2251
R4 2252
R5 2253
Implementation Example 2254
Revised CLI Commands 2254
Glossary 2254
OSPFv3 2256
Enable OSPFv3 on an Interface 2257
Topology 2257
Configuration 2257
Validation 2258
Set Priority 2260
Topology 2260
Configuration 2260
Validation 2261
Area Border Router 2264
Topology 2264
Configuration 2264
Validation 2265
Redistribute Routes into OSPFv3 2269
Topology 2269
Configuration 2269
Validation 2270

Cost 2274
Topology 2274
Configuration 2274
Validation 2276
Virtual Links 2283
Topology 2283
Configuration 2283
Validation 2285
Multiple Instances 2289
Topology 2289
Enable Multiple OSPFv3 Instances on a Router Based on Tags 2289
Configuration 2289
Validation 2291
Redistribute among Multiple Instances 2293
Redistribute with Metric Option 2295
Redistribute with Type Option 2298
Redistribute with Route-Map Option 2300
Not-So-Stubby Area 2303
Topology 2303
NSSA with Route Option 2303
Topology 2304
Configuration 2304
NSSA with the Summary Address Option 2307
Topology 2307
Configuration 2307
NSSA with the Translator Role Option 2311
Topology 2311
Configuration 2311
Link LSA Suppression 2316
Topology 2316
Configuration 2316
Originate Type-7 LSAs and Translate to Type-5 2318
Topology 2318
Configuration 2318
Summarize Inter-Area and External Routes 2322
Topology 2322
Configuration 2322
Validation 2323
Distribute List 2327
Topology 2327
Configuration 2327
OSPFv3 Authentication 2331
Topology 2331
Configuration 2331
Validation 2332
OSPFv3 Graceful Restart Configuration 2336
Topology 2336

Configure R1 for Graceful Restart 2336
Configure R2 as Restart Helper 2337
Remove Capability Restart Configuration from R1 2337
Remove Helper Configuration from R2 2337
Validation 2338
R1 After Configuring Graceful Restart 2338
R1 Before Graceful Restart 2338
R2 Before Graceful Restart 2339
R1 During graceful restart 2340
R2 During graceful restart 2340
R1 After graceful restart 2341
R2 After graceful restart 2341
OSPF Sham-link for VPN Sites Configuration 2342
Topology 2342
CE1 Configuration 2342
PE1 Configuration 2343
PE2 Configuration 2345
CE2 Configuration 2347
Validation 2348
CE1 2348
PE1 2348
PE2 2349
CE2 2350
OSPF TE-Metric Extension 2352
Topology 2352
Configuration 2352
Validation 2353
OSPFv2 LDP Remote Loop-Free Alternate (RLFA) 2358
RLFA Terminology 2358
Repair tunnel 2359
P-space 2359
Extended P-space 2359
Q-space 2359
PQ node 2360
Remote LFA (RLFA) 2360
Establishing RLFA Tunnel 2360
OSPF 2360
LDP 2360
NSM 2360
HSL 2360
Configure LDP Remote Loop-Free Alternate (RLFA) 2360
Assumptions and Limitations 2361
Topology 2361
R1 2361
R2 2363
R3 2364
R4 2366

R5 2367
R6 2368
Validation 2370

Open Shortest Path First Command Reference 2376

OSPFv2 Commands 2382
area authentication 2386
Command Syntax 2386
Parameters 2386
Default 2386
Command Mode 2386
Applicability 2386
Examples 2386
area default-cost 2387
Command Syntax 2387
Parameters 2387
Default 2387
Command Mode 2387
Applicability 2387
Examples 2387
area filter-list 2388
Command Syntax 2388
Parameters 2388
Default 2388
Command Mode 2388
Applicability 2388
Examples 2388
area interface 2390
Command Syntax 2390
Parameters 2390
Default 2390
Command Mode 2390
Applicability 2390
Examples 2390
area interface authentication 2391
Command Syntax 2391
Parameters 2391
Default 2391
Command Mode 2391
Applicability 2391
Examples 2391
area interface cost 2392
Command Syntax 2392
Parameters 2392
Default 2392
Command Mode 2392
Applicability 2392

Examples 2392
area interface database-filter 2393
Command Syntax 2393
Parameters 2393
Default 2393
Command Mode 2393
Applicability 2393
Examples 2393
area interface dead-interval 2394
Command Syntax 2394
Parameters 2394
Default 2394
Command Mode 2394
Applicability 2394
Examples 2394
area interface hello-interval 2395
Command Syntax 2395
Parameters 2395
Default 2395
Command Mode 2395
Applicability 2395
Examples 2395
area interface IFNAME bfd 2396
Command Syntax 2396
Parameters 2396
Default 2396
Command Mode 2396
Applicability 2396
Examples 2396
area interface IFNAME priority 2397
Command Syntax 2397
Parameters 2397
Default 2397
Command Mode 2397
Applicability 2397
Examples 2397
area interface mpls ldp-igp 2398
Command Syntax 2398
Parameters 2398
Default 2398
Command Mode 2398
Applicability 2398
Examples 2398
area interface mpls ldp-igp sync ospf holddown-timer 2399
Command Syntax 2399
Parameters 2399
Default 2399

Command Mode 2399
Applicability 2399
Examples 2399
area interface network-type 2400
Command Syntax 2400
Parameters 2400
Default 2400
Command Mode 2400
Applicability 2400
Examples 2400
area interface passive 2401
Command Syntax 2401
Parameters 2401
Default 2401
Command Mode 2401
Applicability 2401
Examples 2401
area interface retransmit-interval 2402
Command Syntax 2402
Parameters 2402
Default 2402
Command Mode 2402
Applicability 2402
Examples 2402
area interface te-metric 2403
Command Syntax 2403
Parameters 2403
Default 2403
Command Mode 2403
Applicability 2403
Examples 2403
area nssa 2404
Command Syntax 2404
Parameters 2404
Default 2405
Command Mode 2405
Applicability 2405
Examples 2405
area range 2406
Command Syntax 2406
Parameters 2406
Default 2406
Command Mode 2406
Applicability 2406
Examples 2406
area sham-link 2408
Command Syntax 2408

Parameters 2408
Default 2408
Command Mode 2408
Applicability 2408
Examples 2408
area shortcut 2410
Command Syntax 2410
Parameters 2410
Default 2410
Command Mode 2410
Applicability 2410
Examples 2410
area stub 2411
Command Syntax 2411
Parameters 2411
Default 2411
Command Mode 2411
Applicability 2411
Examples 2411
area virtual-link 2412
Command Syntax 2412
Parameters 2412
Default 2413
Command Mode 2413
Applicability 2413
Examples 2413
auto-cost reference bandwidth 2414
Command Syntax 2414
Parameters 2414
Default 2414
Command Mode 2414
Applicability 2414
Examples 2414
bfd all-interfaces 2415
Command Syntax 2415
Parameters 2415
Default 2415
Command Mode 2415
Applicability 2415
Examples 2415
capability cspf 2416
Command Syntax 2416
Parameters 2416
Default 2416
Command Mode 2416
Applicability 2416
Examples 2416

capability lls 2417
Command Syntax 2417
Parameters 2417
Default 2417
Command Mode 2417
Applicability 2417
Examples 2417
capability opaque 2418
Command Syntax 2418
Parameters 2418
Default 2418
Command Mode 2418
Applicability 2418
Examples 2418
capability traffic-engineering 2419
Command Syntax 2419
Parameters 2419
Default 2419
Command Mode 2419
Applicability 2419
Examples 2419
capability vrf-lite 2420
Command Syntax 2420
Parameters 2420
Default 2420
Command Mode 2420
Applicability 2420
Examples 2420
clear ip ospf 2421
Command Syntax 2421
Parameter 2421
Default 2421
Command Mode 2421
Applicability 2421
Examples 2421
compatible rfc1583 2422
Command Syntax 2422
Parameters 2422
Default 2422
Command Mode 2422
Applicability 2422
Examples 2422
debug ip ospf dist-ls 2423
Command Syntax 2423
Parameters 2423
Default 2423
Command Mode 2423

Applicability 2423
Examples 2423
debug ip ospf lfa 2424
Command Syntax 2424
Parameters 2424
Default 2424
Command Mode 2424
Applicability 2424
Examples 2424
debug ip ospf redist 2425
Command Syntax 2425
Parameters 2425
Default 2425
Command Mode 2425
Applicability 2425
Examples 2425
debug ip ospf retransmission 2426
Command Syntax 2426
Parameters 2426
Default 2426
Command Mode 2426
Applicability 2426
Examples 2426
debug ospf 2427
Command Syntax 2427
Parameters 2427
Default 2427
Command Mode 2428
Applicability 2428
Examples 2428
debug ospf database-timer rate-limit 2429
Command Syntax 2429
Parameters 2429
Default 2429
Command Mode 2429
Applicability 2429
Examples 2429
debug ospf events 2430
Command Syntax 2430
Parameters 2430
Default 2430
Command Mode 2430
Applicability 2430
Examples 2430
debug ospf ifsm 2431
Command Syntax 2431
Parameters 2431

Default 2431
Command Mode 2431
Applicability 2431
Examples 2431
debug ospf lsa 2432
Command Syntax 2432
Parameters 2432
Default 2432
Command Mode 2432
Applicability 2432
Examples 2432
debug ospf nfsm 2433
Command Syntax 2433
Parameters 2433
Default 2433
Command Mode 2433
Applicability 2433
Examples 2433
debug ospf nsm 2434
Command Syntax 2434
Parameters 2434
Default 2434
Command Mode 2434
Applicability 2434
Example 2434
debug ospf packet 2435
Command Syntax 2435
Parameters 2435
Default 2435
Command Mode 2435
Applicability 2435
Examples 2435
debug ospf rib 2436
Command Syntax 2436
Parameters 2436
Default 2436
Command Mode 2436
Applicability 2436
Example 2436
debug ospf route 2437
Command Syntax 2437
Parameters 2437
Default 2437
Command Mode 2437
Applicability 2437
Examples 2437
default-information originate 2438

Command Syntax 2438
Parameters 2438
Default 2438
Command Mode 2438
Applicability 2438
Examples 2439
default-metric 2440
Command Syntax 2440
Parameters 2440
Default 2440
Command Mode 2440
Applicability 2440
Examples 2440
distance 2441
Command Syntax 2441
Parameters 2441
Default 2441
Command Mode 2441
Applicability 2441
Examples 2441
distribute bgp-ls 2443
distribute-list 2444
Command Syntax 2444
Parameters 2444
Default 2444
Command Mode 2444
Applicability 2444
Examples 2445
enable db-summary-opt 2446
Command Syntax 2446
Parameters 2446
Default 2446
Command Mode 2446
Applicability 2446
Examples 2446
enable ext-ospf-multi-inst 2447
Command Syntax 2447
Parameters 2447
Default 2447
Command Mode 2447
Applicability 2447
Examples 2447
fast-reroute keep-all-paths 2448
Command Syntax 2448
Parameters 2448
Default 2448
Command Mode 2448

Applicability 2448
Examples 2448
fast-reroute per-prefix remote-lfa area 2449
Command Syntax 2449
Parameters 2449
Default 2449
Command Mode 2449
Applicability 2449
Examples 2449
fast-reroute per-prefix route-map 2450
Command Syntax 2450
Parameters 2450
Default 2450
Command Mode 2450
Applicability 2450
Examples 2450
fast-reroute terminate-hold-on interval 2451
Command Syntax 2451
Parameters 2451
Default 2451
Command Mode 2451
Applicability 2451
Examples 2451
fast-reroute tie-break 2452
Command Syntax 2452
Parameters 2452
Default 2453
Command Mode 2453
Applicability 2453
Examples 2453
host area 2454
Command Syntax 2454
Parameters 2454
Default 2454
Command Mode 2454
Applicability 2454
Examples 2454
ip ospf authentication 2455
Command Syntax 2455
Parameters 2455
Default 2455
Command Mode 2455
Applicability 2455
Examples 2455
ip ospf authentication-key 2456
Command Syntax 2456
Parameters 2456

Default 2456
Command Mode 2456
Applicability 2456
Examples 2456
ip ospf bfd 2458
Command Syntax 2458
Parameters 2458
Default 2458
Command Mode 2458
Applicability 2458
Examples 2458
ip ospf cost 2459
Command Syntax 2459
Parameters 2459
Default 2459
Command Mode 2459
Applicability 2459
Examples 2459
ip ospf database-filter 2460
Command Syntax 2460
Parameters 2460
Default 2460
Command Mode 2460
Applicability 2460
Examples 2460
ip ospf dead-interval 2461
Command Syntax 2461
Parameters 2461
Default 2461
Command Mode 2461
Applicability 2461
Examples 2461
ip ospf demand-circuit 2462
Command Syntax 2462
Parameters 2462
Command Mode 2462
Applicability 2462
Examples 2462
ip ospf disable 2463
Command Syntax 2463
Parameters 2463
Default 2463
Command Mode 2463
Applicability 2463
Example 2463
ip ospf fast-reroute per-prefix candidate disable 2464
Command Syntax 2464

Parameters 2464
Default 2464
Command Mode 2464
Applicability 2464
Example 2464
ip ospf flood-reduction 2465
Command Syntax 2465
Parameters 2465
Default 2465
Command Mode 2465
Applicability 2465
Example 2465
ip ospf hello-interval 2466
Command Syntax 2466
Parameters 2466
Default 2466
Command Mode 2466
Applicability 2466
Examples 2466
ip ospf message-digest-key 2467
Command Syntax 2467
Parameters 2467
Default 2467
Command Mode 2467
Applicability 2467
Examples 2467
ip ospf mtu 2469
Command Syntax 2469
Parameters 2469
Default 2469
Command Mode 2469
Applicability 2469
Examples 2469
ip ospf mtu-ignore 2470
Command Syntax 2470
Parameters 2470
Default 2470
Command Mode 2470
Applicability 2470
Examples 2470
ip ospf multi-area 2471
Command Syntax 2471
Parameters 2471
Default 2471
Command Mode 2471
Applicability 2471
Examples 2471

ip ospf network 2473
Command Syntax 2473
Parameters 2473
Default 2473
Command Mode 2473
Applicability 2473
Examples 2473
ip ospf priority 2474
Command Syntax 2474
Parameters 2474
Default 2474
Command Mode 2474
Applicability 2474
Examples 2474
ip ospf resync-timeout 2475
Command Syntax 2475
Parameters 2475
Default 2475
Command Mode 2475
Applicability 2475
Examples 2475
ip ospf retransmit-interval 2476
Command Syntax 2476
Parameters 2476
Default 2476
Command Mode 2476
Applicability 2476
Examples 2476
ip ospf transmit-delay 2477
Command Syntax 2477
Parameters 2477
Default 2477
Command Mode 2477
Applicability 2477
Examples 2477
log-adjacency-changes 2478
Command Syntax 2478
Parameters 2478
Default 2478
Command Mode 2478
Applicability 2478
Examples 2478
max-concurrent-dd 2479
Command Syntax 2479
Parameters 2479
Default 2479
Command Mode 2479

Applicability 2479
Examples 2479
maximum-area 2480
Command Syntax 2480
Parameters 2480
Default 2480
Command Mode 2480
Applicability 2480
Examples 2480
max-metric 2481
Command Syntax 2481
Parameters 2481
Default 2481
Command Mode 2481
Applicability 2481
Examples 2481
neighbor 2483
Command Syntax 2483
Parameters 2483
Default 2483
Command Mode 2483
Applicability 2483
Examples 2483
network 2484
Command Syntax 2484
Parameters 2484
Default 2484
Command Mode 2484
Applicability 2484
Examples 2484
ospf abr-type 2486
Command Syntax 2486
Parameters 2486
Default 2486
Command Mode 2486
Applicability 2486
Examples 2486
ospf area-interface-config-mode 2487
Command Syntax 2487
Parameters 2487
Default 2487
Command Mode 2487
Applicability 2487
Examples 2487
ospf flood-reduction 2488
Command Syntax 2488
Parameters 2488

Default 2488
Command Mode 2488
Applicability 2488
Examples 2488
ospf point-point rfc-incompatible 2489
Command Syntax 2489
Parameters 2489
Default 2489
Command Mode 2489
Applicability 2489
Examples 2489
ospf router-id 2490
Command Syntax 2490
Parameters 2490
Default 2490
Command Mode 2490
Applicability 2490
Examples 2490
overflow database 2491
Command Syntax 2491
Parameters 2491
Default 2491
Command Mode 2491
Applicability 2491
Examples 2491
overflow database external 2492
Command Syntax 2492
Parameters 2492
Default 2492
Command Mode 2492
Applicability 2492
Examples 2492
passive-interface 2493
Command Syntax 2493
Parameters 2493
Default 2493
Command Mode 2493
Applicability 2493
Examples 2493
redistribute 2495
Command Syntax 2495
Parameters 2495
Default 2495
Command Mode 2496
Applicability 2496
Examples 2496
router ospf 2497

Command Syntax 2497
Parameters 2497
Default 2497
Command Mode 2497
Applicability 2497
Examples 2497
show cspf lsp 2497
Command Syntax 2498
Parameters 2498
Command Mode 2498
Applicability 2498
Example 2498
show cspf rsvp forwarding-timer 2500
Command Syntax 2500
Parameters 2500
Default 2500
Command Mode 2500
Applicability 2500
Example 2500
show debugging ospf 2501
Command Syntax 2501
Parameters 2501
Default 2501
Command Mode 2501
Applicability 2501
Example 2501
show ip ospf 2502
Command Syntax 2502
Parameters 2502
Default 2502
Command Mode 2502
Applicability 2502
Examples 2502
show ip ospf border-routers 2506
Command Syntax 2506
Parameters 2506
Default 2506
Command Mode 2506
Applicability 2506
Example 2506
show ip ospf database brief 2508
Command Syntax 2508
Parameters 2508
Default 2508
Command Mode 2508
Applicability 2508
Examples 2508

© 2026 IP Infusion Inc. Proprietary

show ip ospf database detail 2510
Command Syntax 2510
Parameters 2510
Default 2510
Command Mode 2511
Applicability 2511
Examples 2511
show ip ospf igp-shortcut-lsp 2518
Command Syntax 2518
Parameters 2518
Default 2518
Command Mode 2518
Applicability 2518
Example 2518
show ip ospf igp-shortcut-route 2519
Command Syntax 2519
Parameters 2519
Default 2519
Command Mode 2519
Applicability 2519
Example 2519
show ip ospf interface 2520
Command Syntax 2520
Parameters 2520
Default 2520
Command Mode 2520
Applicability 2520
Example 2520
show ip ospf multi-area-adjacencies 2523
Command Syntax 2523
Parameters 2523
Default 2523
Command Mode 2523
Applicability 2523
Example 2523
show ip ospf neighbor 2525
Command Syntax 2525
Parameters 2525
Default 2525
Command Mode 2525
Applicability 2525
Example 2525
show ip ospf route 2529
Command Syntax 2529
Parameters 2529
Default 2529
Command Mode 2529

Applicability 2529
Example 2529
show ip ospf sham-links 2532
Command Syntax 2532
Parameters 2532
Default 2532
Command Mode 2532
Applicability 2532
Examples 2532
show ip ospf valid 2534
Command Syntax 2534
Parameters 2534
Default 2534
Command Mode 2534
Applicability 2534
Examples 2534
show ip ospf virtual-links 2535
Command Syntax 2535
Parameters 2535
Default 2535
Command Mode 2535
Applicability 2535
Examples 2535
show ip protocols 2537
Command Syntax 2537
Parameters 2537
Default 2537
Command Mode 2537
Applicability 2537
Examples 2537
show ip route fast-reroute 2539
Command Syntax 2539
Parameters 2539
Default 2539
Command Mode 2539
Applicability 2539
Examples 2539
shutdown 2540
Command Syntax 2540
Parameters 2540
Default 2540
Command Mode 2540
Applicability 2540
Examples 2540
snmp context-name 2541
Command Syntax 2541
Parameters 2541

Default 2541
Command Mode 2541
Applicability 2541
Examples 2541
snmp restart ospf 2542
Command Syntax 2542
Parameter 2542
Default 2542
Command Mode 2542
Applicability 2542
Examples 2542
summary-address 2543
Command Syntax 2543
Parameters 2543
Default 2543
Command Mode 2543
Applicability 2543
Examples 2543
te-metric 2545
Command Syntax 2545
Parameters 2545
Default 2545
Command Mode 2545
Applicability 2545
Examples 2545
timers lsa arrival 2546
Command Syntax 2546
Parameters 2546
Default 2546
Command Mode 2546
Applicability 2546
Examples 2546
timers spf exp 2547
Command Syntax 2547
Parameters 2547
Default 2547
Command Mode 2547
Applicability 2547
Examples 2547
timers throttle lsa 2548
Command Syntax 2548
Parameters 2548
Default 2548
Command Mode 2548
Applicability 2548
Examples 2548
OSPFv2 Graceful Restart Commands 2549

action interface 2550
Command Syntax 2550
Parameters 2550
Command Mode 2550
Default 2550
Applicability 2550
Example 2550
capability restart 2551
Command Syntax 2551
Parameters 2551
Default 2551
Command Mode 2551
Applicability 2551
Examples 2551
event nbr 2552
Command Syntax 2552
Parameters 2552
Default 2552
Command Mode 2552
Applicability 2552
Example 2552
debug ip ospf graceful-restart 2553
Command Syntax 2553
Parameters 2553
Command Mode 2553
Applicability 2553
Examples 2553
if-event-track ospf event 2554
Command Syntax 2554
Parameters 2554
Default 2554
Command Mode 2554
Applicability 2554
Example 2554
ospf restart grace-period 2555
Command Syntax 2555
Parameters 2555
Default 2555
Command Mode 2555
Applicability 2555
Examples 2555
ospf restart helper 2556
Command Syntax 2556
Parameters 2556
Default 2556
Command Mode 2556
Applicability 2556

Examples 2556
restart ip ospf graceful 2557
Command Syntax 2557
Parameters 2557
Default 2557
Command Mode 2557
Applicability 2557
Examples 2557
OSPFv3 Graceful Restart Commands 2558
capability restart 2559
Command Syntax 2559
Parameter 2559
Default 2559
Command Mode 2559
Applicability 2559
Examples 2559
ipv6 ospf restart grace-period 2560
Command Syntax 2560
Parameters 2560
Default 2560
Command Mode 2560
Applicability 2560
Examples 2560
ipv6 ospf restart helper 2561
Command Syntax 2561
Parameters 2561
Default 2561
Command Mode 2561
Applicability 2561
Examples 2561
ipv6 ospf restart planned-only 2562
Command Syntax 2562
Parameters 2562
Default 2562
Command Mode 2562
Applicability 2562
Example 2562
restart ipv6 ospf graceful 2563
Command Syntax 2563
Parameters 2563
Default 2563
Command Mode 2563
Applicability 2563
Examples 2563
OSPFv3 Commands 2564
abr-type 2566
Command Syntax 2566

Parameters 2566
Default 2566
Command Mode 2566
Applicability 2566
Examples 2566
area default-cost 2567
Command Syntax 2567
Parameters 2567
Default 2567
Command Mode 2567
Applicability 2567
Examples 2567
area nssa 2568
Command Syntax 2568
Parameters 2568
Default 2569
Command Mode 2569
Applicability 2569
Examples 2569
area range 2570
Command Syntax 2570
Parameters 2570
Default 2570
Command Mode 2570
Applicability 2570
Examples 2570
area stub 2571
Command Syntax 2571
Parameters 2571
Default 2571
Command Mode 2571
Applicability 2571
Examples 2571
area virtual-link 2572
Command Syntax 2572
Parameters 2572
Default 2573
Command Mode 2573
Applicability 2573
Examples 2573
auto-cost reference bandwidth 2574
Command Syntax 2574
Parameters 2574
Default 2574
Command Mode 2574
Applicability 2574
Examples 2574

© 2026 IP Infusion Inc. Proprietary

bfd all-interfaces 2575
Command Syntax 2575
Parameters 2575
Default 2575
Command Mode 2575
Applicability 2575
Examples 2575
clear ipv6 ospf process 2576
Command Syntax 2576
Parameters 2576
Command Mode 2576
Applicability 2576
Examples 2576
debug ipv6 ospf 2577
Command Syntax 2577
Parameters 2577
Command Mode 2577
Applicability 2577
Examples 2577
debug ipv6 ospf bfd 2578
Command Syntax 2578
Parameters 2578
Command Mode 2578
Applicability 2578
Examples 2578
debug ipv6 ospf events 2579
Command Syntax 2579
Parameters 2579
Command Mode 2579
Applicability 2579
Examples 2579
debug ipv6 ospf ifsm 2580
Command Syntax 2580
Parameters 2580
Command Mode 2580
Applicability 2580
Examples 2580
debug ipv6 ospf lsa 2581
Command Syntax 2581
Parameters 2581
Command Mode 2581
Applicability 2581
Examples 2581
debug ipv6 ospf nfsm 2582
Command Syntax 2582
Parameters 2582
Command Mode 2582

© 2026 IP Infusion Inc. Proprietary

Applicability 2582
Examples 2582
debug ipv6 ospf nsm 2583
Command Syntax 2583
Parameters 2583
Command Mode 2583
Applicability 2583
Examples 2583
debug ipv6 ospf packet 2584
Command Syntax 2584
Parameters 2584
Command Mode 2584
Applicability 2584
Examples 2584
debug ipv6 ospf retransmission 2585
Command Syntax 2585
Parameters 2585
Command Mode 2585
Applicability 2585
Examples 2585
debug ipv6 ospf rib 2586
Command Syntax 2586
Parameters 2586
Command Mode 2586
Applicability 2586
Examples 2586
debug ipv6 ospf route 2587
Command Syntax 2587
Parameters 2587
Command Mode 2587
Applicability 2587
Examples 2587
default-information originate 2588
Command Syntax 2588
Parameters 2588
Default 2588
Command Mode 2588
Applicability 2589
Examples 2589
default-metric 2590
Command Syntax 2590
Parameter 2590
Default 2590
Command Mode 2590
Applicability 2590
Examples 2590
distance 2591

Command Syntax 2591
Parameters 2591
Command Mode 2591
Defaults 2591
Applicability 2591
Examples 2591
distribute bgp-ls 2592
Command Syntax 2592
Parameters 2592
Default 2592
Command Mode 2592
Applicability 2592
Examples 2592
distribute-list 2593
Command Syntax 2593
Parameters 2593
Default 2593
Command Mode 2593
Applicability 2594
Examples 2594
enable db-summary-opt 2595
Command Syntax 2595
Parameters 2595
Default 2595
Command Mode 2595
Applicability 2595
Examples 2595
ipv6 ospf authentication 2596
Command Syntax 2596
Parameters 2596
Default 2596
Command Mode 2596
Applicability 2596
Example 2596
ipv6 ospf bfd 2597
Command Syntax 2597
Parameters 2597
Default 2597
Command Mode 2597
Applicability 2597
Example 2597
ipv6 ospf cost 2598
Command Syntax 2598
Parameters 2598
Default 2598
Command Mode 2598
Applicability 2598

Examples 2598
ipv6 ospf dead-interval 2599
Command Syntax 2599
Parameters 2599
Default 2599
Command Mode 2599
Applicability 2599
Examples 2599
ipv6 ospf display route single-line 2600
Command Syntax 2600
Parameters 2600
Default 2600
Command Mode 2600
Applicability 2600
Examples 2600
ipv6 ospf hello-interval 2601
Command Syntax 2601
Parameters 2601
Default 2601
Command Mode 2601
Applicability 2601
Examples 2601
ipv6 ospf link-lsa-suppression 2602
Command Syntax 2602
Parameters 2602
Default 2602
Command Mode 2602
Applicability 2602
Examples 2602
ipv6 ospf mtu 2603
Command syntax 2603
Parameters 2603
Default 2603
Command Mode 2603
Applicability 2603
Examples 2603
ipv6 ospf mtu-ignore 2604
Command syntax 2604
Parameters 2604
Default 2604
Command Mode 2604
Applicability 2604
Examples 2604
ipv6 ospf neighbor 2605
Command Syntax 2605
Parameters 2605
Default 2605

Command Mode 2606
Applicability 2606
Examples 2606
ipv6 ospf network 2607
Command Syntax 2607
Parameters 2607
Default 2607
Command Mode 2607
Applicability 2607
Examples 2607
ipv6 ospf priority 2608
Command Syntax 2608
Parameters 2608
Default 2608
Applicability 2608
Examples 2608
ipv6 ospf retransmit-interval 2609
Command Syntax 2609
Parameters 2609
Default 2609
Command Mode 2609
Applicability 2609
Examples 2609
ipv6 ospf transmit-delay 2610
Command Syntax 2610
Parameters 2610
Default 2610
Command Mode 2610
Applicability 2610
Examples 2610
ipv6 router ospf 2611
Command Syntax 2611
Parameters 2611
Default 2611
Command Mode 2611
Applicability 2612
Examples 2612
ipv6 te-metric 2613
Command Syntax 2613
Parameters 2613
Default 2613
Command Mode 2613
Applicability 2613
Examples 2613
log-adjacency-changes 2614
Command Syntax 2614
Parameters 2614

Default 2614
Command Mode 2614
Applicability 2614
Example 2614
max-concurrent-dd 2615
Command Syntax 2615
Parameters 2615
Default 2615
Command Mode 2615
Applicability 2615
Examples 2615
passive-interface 2616
Command Syntax 2616
Parameters 2616
Default 2616
Command Mode 2616
Applicability 2616
Examples 2616
redistribute 2617
Command Syntax 2617
Parameters 2617
Default 2618
Command Mode 2618
Applicability 2618
Examples 2618
router-id 2619
Command Syntax 2619
Parameters 2619
Default 2619
Command Mode 2619
Applicability 2619
Examples 2619
router ipv6 ospf 2620
Command Syntax 2620
Parameters 2620
Default 2620
Command Mode 2620
Applicability 2620
Examples 2620
show debugging ipv6 ospf 2621
Command Syntax 2621
Parameters 2621
Command Mode 2621
Applicability 2621
Examples 2621
show ipv6 ospf 2622
Command Syntax 2622

Parameters 2622
Command Mode 2622
Applicability 2622
Example 2622
show ipv6 ospf database 2624
Command Syntax 2624
Parameters 2624
Command Mode 2625
Applicability 2625
Example: adv-router Parameter 2625
Example: intra-prefix and adv-router Parameters 2627
show ipv6 ospf interface 2628
Command Syntax 2628
Parameters 2628
Command Mode 2628
Applicability 2628
Usage 2628
show ipv6 ospf neighbor 2630
Command Syntax 2630
Parameters 2630
Command Mode 2630
Applicability 2630
Example 2630
show ipv6 ospf route 2633
Command Syntax 2633
Parameters 2633
Command Mode 2633
Applicability 2633
Example 2633
show ipv6 ospf virtual-links 2635
Command Syntax 2635
Parameters 2635
Command Mode 2635
Applicability 2635
Example 2635
show ipv6 ospfv3 topology 2637
Command Syntax 2637
Parameters 2637
Command Mode 2637
Applicability 2637
Example 2637
Example 2637
show ipv6 route fast-reroute 2639
Command Syntax 2639
Parameters 2639
Command Mode 2639
Applicability 2639

Usage 2639
show ipv6 vrf 2640
Command Syntax 2640
Parameter 2640
Command Mode 2640
Applicability 2640
Examples 2640
shutdown 2641
Command Syntax 2641
Parameter 2641
Default 2641
Command Mode 2641
Applicability 2641
Example 2641
snmp restart ospf6 2642
Command Syntax 2642
Parameter 2642
Default 2642
Command Mode 2642
Applicability 2642
Examples 2642
summary-address 2643
Command Syntax 2643
Parameters 2643
Default 2643
Command Mode 2643
Applicability 2644
Examples 2644
timers spf exp 2645
Command Syntax 2645
Parameters 2645
Default 2645
Command Mode 2645
Applicability 2645
Example 2645
OSPF VPN Commands 2646
capability vrf-lite 2647
Command Syntax 2647
Parameters 2647
Default 2647
Command Mode 2647
Applicability 2647
Example 2647
domain-id 2648
Command Syntax 2648
Parameters 2648
Default 2648

Command Mode 2648
Applicability 2648
Example 2649
router ospf vrf 2650
Command Syntax 2650
Parameters 2650
Default 2650
Command Mode 2650
Applicability 2650
Examples 2650
CSPF-TE Commands 2651
cspf default-retry-interval 2652
Command Syntax 2652
Parameter 2652
Default 2652
Command Mode 2652
Applicability 2652
Examples 2652
cspf disable-better-protection 2653
Command Syntax 2653
Parameter 2653
Default 2653
Command Mode 2653
Applicability 2653
Examples 2653
cspf enable-better-protection 2654
Command Syntax 2654
Parameter 2654
Default 2654
Command Mode 2654
Applicability 2654
Examples 2654
cspf enable-resource-constraint-crankback 2655
Command Syntax 2655
Parameter 2655
Default 2655
Command Mode 2655
Applicability 2655
Examples 2655
cspf tie-break 2656
Command Syntax 2656
Parameters 2656
Default 2656
Command Mode 2656
Applicability 2656
Examples 2656

Routing Information Protocol Configuration 2657
Router Information Protocol Configuration 2658
Enable RIP 2658
Topology 2658
Validation 2658
Specify RIP Version 2661
Topology 2661
Validation 2662
Authentication with a Single Key 2664
Topology 2664
Validation 2665
Text Authentication with Multiple Keys 2669
Topology 2669
Validation 2673
R2 2675
MD5 Authentication with Multiple Keys 2677
Topology 2677
Validation 2681
RIPV2 VRF Configuration 2683
Topology 2684
RIPV2 VRF Configuration 2684
Validation 2685
RIPng 2686
Topology 2686
Configuration 2686
R1 2686
R2 2686
Validation 2687
R1 2687
R2 2688

Routing Information Protocol Command Reference 2689
Routing Information Protocol Commands 2691
accept-lifetime 2693
Command Syntax 2693
Parameter 2693
Default 2693
Command Mode 2694
Applicability 2694
Examples 2694
cisco-metric-behavior 2695
Command Syntax 2695
Parameters 2695
Default 2695
Command Mode 2695
Applicability 2695
Example 2695

© 2026 IP Infusion Inc. Proprietary

clear ip rip route 2696
Command Syntax 2696
Parameters 2696
Default 2696
Command Mode 2696
Applicability 2696
Examples 2697
clear ip rip route vrf NAME 2698
Command Syntax 2698
Parameters 2698
Default 2698
Command Mode 2698
Applicability 2698
Examples 2698
clear ip rip statistics 2699
Command Syntax 2699
Parameters 2699
Default 2699
Command Mode 2699
Applicability 2699
Examples 2699
debug rip 2700
Command Syntax 2700
Parameters 2700
Default 2700
Command Mode 2700
Applicability 2700
Examples 2700
default-information originate 2702
Command Syntax 2702
Parameters 2702
Default 2702
Command Mode* 2702
Applicability 2702
Examples 2702
default-metric 2703
Command Syntax 2703
Parameter 2703
Default 2703
Command Mode 2703
Applicability 2703
Examples 2703
distance 2704
Command Syntax 2704
Parameters 2704
Default 2704
Command Mode 2704

© 2026 IP Infusion Inc. Proprietary

Applicability 2704
Examples 2704
distribute-list 2705
Command Syntax 2705
Parameters 2705
Default 2705
Command Mode 2705
Applicability 2705
Example 2705
ip rip authentication key-chain 2706
Command Syntax 2706
Parameters 2706
Default 2706
Command Mode 2706
Applicability 2706
Examples 2706
ip rip authentication mode 2707
Command Syntax 2707
Parameters 2707
Default 2707
Command Mode 2707
Applicability 2707
Examples 2707
ip rip authentication string 2708
Command Syntax 2708
Parameters 2708
Default 2708
Command Mode 2708
Applicability 2708
Examples 2708
ip rip receive-packet 2709
Command Syntax 2709
Parameters 2709
Default 2709
Command Mode 2709
Applicability 2709
Example 2709
ip rip receive version 2710
Command Syntax 2710
Parameters 2710
Default 2710
Command Mode 2710
Applicability 2710
Examples 2710
ip rip send-packet 2711
Command Syntax 2711
Parameters 2711

Default 2711
Command Mode 2711
Applicability 2711
Example 2711
ip rip send version 2712
Command Syntax 2712
Parameters 2712
Default 2712
Command Mode 2712
Applicability 2712
Examples 2712
ip rip split-horizon 2713
Command Syntax 2713
Parameter 2713
Default 2713
Command Mode 2713
Applicability 2713
Examples 2713
key chain 2714
Command Syntax 2714
Parameters 2714
Default 2714
Command Mode 2714
Applicability 2714
Examples 2714
key-id 2715
Command Syntax 2715
Parameters 2715
Default 2715
Command Mode 2715
Applicability 2715
Examples 2715
key-string 2716
Command Syntax 2716
Parameters 2716
Default 2716
Command Mode 2716
Applicability 2716
Examples 2716
key-string encrypted 2717
Command Syntax 2717
Parameters 2717
Default 2717
Command Mode 2717
Applicability 2717
Example 2717
maximum-prefix 2718

Command Syntax 2718
Parameters 2718
Default 2718
Command Mode 2718
Applicability 2718
Examples 2718
neighbor 2719
Command Syntax 2719
Parameter 2719
Default 2719
Command Mode 2719
Applicability 2719
Examples 2719
network 2720
Command Syntax 2720
Parameters 2720
Default 2720
Command Mode 2720
Applicability 2720
Examples 2720
offset-list 2721
Command Syntax 2721
Parameters 2721
Default 2721
Command Mode 2721
Applicability 2721
Examples 2721
passive-interface 2722
Command Syntax 2722
Parameters 2722
Default 2722
Command Mode 2722
Applicability 2722
Examples 2722
recv-buffer-size 2723
Command Syntax 2723
Parameters 2723
Default 2723
Command Mode 2723
Applicability 2723
Examples 2723
redistribute 2724
Command Syntax 2724
Parameters 2724
Default 2724
Command Mode 2724
Examples 2725

route 2726
Command Syntax 2726
Parameter 2726
Default 2726
Command Mode 2726
Applicability 2726
Examples 2726
router rip 2727
Command Syntax 2727
Parameter 2727
Default 2727
Command Mode 2727
Applicability 2727
Examples 2727
send-lifetime 2728
Command Syntax 2728
Parameters 2728
Default 2728
Command Mode 2728
Applicability 2729
Examples 2729
show debugging rip 2730
Command Syntax 2730
Parameters 2730
Default 2730
Command Mode 2730
Applicability 2730
Example 2730
show ip protocols rip 2731
Command Syntax 2731
Parameters 2731
Default 2731
Command Mode 2731
Applicability 2731
Example 2731
show ip rip 2733
Command Syntax 2733
Parameters 2733
Default 2733
Command Mode 2733
Applicability 2733
Example 2733
show ip rip interface 2735
Command Syntax 2735
Parameters 2735
Default 2735
Command Mode 2735

Applicability 2735
Example 2735
show ip rip statistics 2737
Command Syntax 2737
Parameters 2737
Default 2737
Command Mode 2737
Applicability 2737
Example 2737
snmp restart rip 2739
Command Syntax 2739
Parameters 2739
Default 2739
Command Mode 2739
Applicability 2739
Examples 2739
timers basic 2740
Command Syntax 2740
Parameters 2740
Default 2740
Command Mode 2740
Applicability 2740
Examples 2740
version 2741
Command Syntax 2741
Parameters 2741
Default 2741
Command Mode 2741
Applicability 2741
Examples 2741
RIPng Commands 2742
aggregate-address 2743
Command Syntax 2743
Parameter 2743
Command Mode 2743
Applicability 2743
Examples 2743
cisco-metric-behavior 2744
Command Syntax 2744
Parameters 2744
Default 2744
Command Mode 2744
Applicability 2744
Example 2744
clear ipv6 rip route 2745
Command Syntax 2745
Parameters 2745

Command Mode 2745
Applicability 2745
Example 2745
debug ipv6 rip 2746
Command Syntax 2746
Parameters 2746
Default 2746
Command Mode 2746
Applicability 2746
Examples 2746
default-information originate 2747
Command Syntax 2747
Parameters 2747
Default 2747
Command Mode 2747
Applicability 2747
Examples 2747
default-metric 2748
Command Syntax 2748
Parameter 2748
Default 2748
Command Mode 2748
Applicability 2748
Examples 2748
distance 2749
Command Syntax 2749
Parameter 2749
Default 2749
Command Mode 2749
Applicability 2749
Example 2749
distribute-list 2750
Command Syntax 2750
Parameters 2750
Default 2750
Command Mode 2750
Applicability 2750
Example 2750
ipv6 rip metric-offset 2751
Command Syntax 2751
Parameter 2751
Default 2751
Command Mode 2751
Applicability 2751
Examples 2751
ipv6 rip split-horizon 2752
Command Syntax 2752

© 2026 IP Infusion Inc. Proprietary

Parameter 2752
Default 2752
Command Mode 2752
Applicability 2752
Examples 2752
ipv6 router rip 2753
Command Syntax 2753
Parameters 2753
Default 2753
Command Mode 2753
Applicability 2753
Example 2753
neighbor 2754
Command Syntax 2754
Parameters 2754
Default 2754
Command Mode 2754
Applicability 2754
Example 2754
offset-list 2755
Command Syntax 2755
Parameters 2755
Default 2755
Command Mode 2755
Applicability 2755
Examples 2755
passive-interface 2756
Command Syntax 2756
Parameters 2756
Default 2756
Command Mode 2756
Applicability 2756
Examples 2756
recv-buffer-size 2757
Command Syntax 2757
Parameters 2757
Default 2757
Command Mode 2757
Applicability 2757
Examples 2757
redistribute 2758
Command Syntax 2758
Parameters 2758
Command Mode 2758
Applicability 2758
Examples 2758
route 2760

© 2026 IP Infusion Inc. Proprietary

Command Syntax 2760
Parameter 2760
Command Mode 2760
Applicability 2760
Examples 2760
route-map 2761
Command Syntax 2761
Parameters 2761
Command Mode 2761
Applicability 2761
Examples 2761
router ipv6 rip 2762
Command Syntax 2762
Parameters 2762
Command Mode 2762
Applicability 2762
Examples 2762
show debugging ipv6 rip 2763
Command Syntax 2763
Parameters 2763
Command Mode 2763
Applicability 2763
Example 2763
show ipv6 protocols rip 2764
Command Syntax 2764
Parameters 2764
Command Mode 2764
Applicability 2764
Example 2764
show ipv6 rip 2765
Command Syntax 2765
Parameters 2765
Command Mode 2765
Applicability 2765
Example 2765
show ipv6 rip interface 2766
Command Syntax 2766
Parameters 2766
Command Mode 2766
Applicability 2766
Example 2766
timers basic 2767
Command Syntax 2767
Parameters 2767
Command Mode 2767
Applicability 2767
Examples 2767

Routing Information Protocol VPN Commands 2768
show ip rip interface vrf 2769
Command Syntax 2769
Parameters 2769
Default 2769
Command Mode 2769
Applicability 2769
Example 2769
show ip rip vrf 2771
Command Syntax 2771
Parameters 2771
Default 2771
Command Mode 2771
Applicability 2771
Examples 2771
show ip vrf 2773
Command Syntax 2773
Parameter 2773
Command Mode 2773
Applicability 2773
Example 2773
Routing Information Protocol Authentication 2774
Single Key Authentication 2774
Example 2774
Multiple Keys Authentication 2774

Two-Way Active Measurement Protocol Configuration 2776

Two-way Active Measurement Protocol 2777
Overview 2777
Periodic Advertisement 2777
Accelerated Advertisement 2777
Limitation and Assumptions 2777
Topology 2778
Configuration 2778
Configuration TWAMP Measurement on ROUTER 2778
Configuring TWAMP Delay Profiles 2779
Configuring Loss Measurement 2780
Un-Configuring/Stopping TWAMP Measurement 2780
Un-Configuring TWAMP Delay profile 2780
Clearing TWAMP Statistics 2781
Validation 2781
End-to-End Measurement 2782
Topology 2783
Configuration 2783
Validation 2785
TWAMP for Client and Server with SNMP 2787
Overview 2787

TWAMP Client 2787
Client Feature Characteristics 2787
TWAMP Server 2787
Server Feature Characteristics 2787
SNMP for TWAMP 2788
SNMP Feature Characteristics 2788
Benefits 2788
Prerequisites 2788
Topology 2788
Server Configuration 2789
Configuration to Trigger TWAMP Traps 2790
Validation 2796
Glossary 2800
TWAMP for VPN Inter-Access Interface 2801
Overview 2801
Feature Characteristics 2801
Benefits 2801
Prerequisites 2801
Limitations 2801
Topology 2802
Configurations 2802
Validation 2806
CLI Commands 2809
twamp session 2810
no twamp session all 2811
shut 2812
no shut 2813
reflector-ip 2814
reflector-port 2815
sender-ip 2816
sender-port 2817
dscp 2818
vrf-name 2819
vpn-net-interface 2820
description 2821
burst-interval 2822
burst-count 2823
interval 2824
notification enabled 2825
notification threshold 2826
notification minimum-change 2827
show twamp-sessions 2828
show twamp-sessions WORD 2829
show twamp-sessions all 2832
show twamp packet-rate 2835
show twamp session count 2836
Glossary 2838

© 2026 IP Infusion Inc. Proprietary

Two-Way Active Measurement Protocol Command Reference 2839
TWAMP Commands 2840
control-admin-state 2842
Command Syntax 2842
Parameters 2842
Default 2842
Command Mode 2842
Applicability 2842
Example 2842
reflector-admin-state 2843
Command Syntax 2843
Parameters 2843
Default 2843
Command Mode 2843
Applicability 2843
Example 2843
twamp stop-test-session (Deprecated) 2844
Command Syntax 2844
Parameters 2844
Default 2844
Command Mode 2844
Applicability 2844
Examples 2844
reflector-name 2845
Command Syntax 2845
Parameters 2845
Default 2845
Command Mode 2845
Applicability 2845
Example 2846
show debugging twamp 2847
Command Syntax 2847
Parameters 2847
Default 2847
Config Mode 2847
Applicability 2847
Examples 2847
show running-config twamp 2848
Command Syntax 2848
Parameters 2848
Command Mode 2848
Applicability 2848
Examples 2848
Show twamp server connections information 2849
Command Syntax 2849
Parameters 2849
Default 2849

Command mode 2849
Applicability 2849
Examples 2849
show twamp server sessions information 2850
Command Syntax 2850
Parameters 2850
Default 2850
Command mode 2850
Applicability 2850
Examples 2850
show twamp server status 2851
Command Syntax 2851
Parameters 2851
Default 2851
Command mode 2851
Applicability 2851
Examples 2851
show twamp statistics 2852
Command Syntax 2852
Parameters 2852
Command Mode 2852
Applicability 2852
Examples 2852
twamp client 2854
Command Syntax 2854
Parameters 2854
Default 2854
Command Mode 2854
Applicability 2854
Examples 2854
twamp client reset connection 2855
Command Syntax 2855
Parameters 2855
Default 2855
Command mode 2855
Applicability 2855
Examples 2855
twamp client start connection 2856
Command Syntax 2856
Parameters 2856
Default 2856
Command mode 2856
Applicability 2856
Examples 2856
twamp client stop connection 2857
Command Syntax 2857
Parameters 2857

Default 2857
Command mode 2857
Applicability 2857
Examples 2857
twamp connection 2858
Command Syntax 2858
Parameters 2858
Default 2858
Command mode 2858
Applicability 2858
Examples 2858
twamp server 2859
Command Syntax 2859
Parameters 2859
Default 2859
Command mode 2859
Applicability 2859
Examples 2859
twamp server inactivity-timeout 2860
Command Syntax 2860
Parameters 2860
Default 2860
Command mode 2860
Applicability 2860
Examples 2860
twamp server listen-vrf 2861
Command Syntax 2861
Parameters 2861
Default 2861
Command mode 2861
Applicability 2861
Examples 2861
twamp server max-connection-duration 2862
Command Syntax 2862
Parameters 2862
Default 2862
Command mode 2862
Applicability 2862
Examples 2862
twamp server max-connections 2863
Command Syntax 2863
Parameters 2863
Default 2863
Command mode 2863
Applicability 2863
Examples 2863
twamp server maximum-sessions-per-connection 2864

Command Syntax 2864
Parameters 2864
Default 2864
Command mode 2864
Applicability 2864
Examples 2864
twamp client maximum-connections 2865
Command Syntax 2865
Parameters 2865
Default 2865
Command mode 2865
Applicability 2865
Examples 2865
twamp client maximum-sessions 2866
Command Syntax 2866
Parameters 2866
Default 2866
Command mode 2866
Applicability 2866
Examples 2866
twamp client maximum-sessions-per-connection 2867
Command Syntax 2867
Parameters 2867
Default 2867
Command mode 2867
Applicability 2867
Examples 2867
twamp client server 2868
Command Syntax 2868
Parameters 2868
Default 2868
Command mode 2868
Applicability 2868
Examples 2868
twamp client session 2869
Command Syntax 2869
Parameters 2869
Default 2869
Command mode 2869
Applicability 2869
Examples 2869
twamp-light control 2870
Command Syntax 2870
Parameters 2870
Command Mode 2870
Applicability 2870
Example 2870

twamp-light reflector 2871
Command Syntax 2871
Parameters 2871
Command Mode 2871
Applicability 2871
Examples 2871
twamp server admin state 2872
Command Syntax 2872
Parameters 2872
Default 2872
Command mode 2872
Applicability 2872
Examples 2872
twamp server max-connections-per-client 2873
Command Syntax 2873
Parameters 2873
Default 2873
Command mode 2873
Applicability 2873
Examples 2873
twamp server maximum-sessions 2874
Command Syntax 2874
Parameters 2874
Default 2874
Command mode 2874
Applicability 2874
Examples 2874
twamp server port 2875
Command Syntax 2875
Parameters 2875
Default 2875
Command mode 2875
Applicability 2875
Examples 2875
twamp start-test-session (Deprecated) 2876
Command Syntax 2876
Parameters 2876
Default 2876
Command Mode 2876
Applicability 2876
Examples 2876
test-session-name (Deprecated) 2877
Command Syntax 2877
Parameters 2877
Default 2877
Command Mode 2877
Applicability 2877

Examples 2878

VRF Lite Configuration 2879

RIP Configuration 2881
Topology 2881
Configuration IPv4 VRF 2881
R1 2881
R1 2882
Validation 2882
Configuration IPv6 VRF 2883
R1 2883
R2 2884
Validation 2885
VRF Configuration 2886
Overview 2886
Topology 2886
Default VRF 2886
Adding a Static Route 2886
User-Defined VRF 2886
Adding a Static Route 2887
Validation 2887
BGP VRF Configuration 2889
Overview 2889
Configurations 2889
Topology 2889
R1 2889
R2 2890
Validation 2890
R1 2890
R2 2891
Static Inter-VRF Route Leaking Configuration 2893
Overview 2893
Topology 2893
IPv4 Inter-VRF Route Leaking Configuration 2893
R1 Configuration 2893
R2 Configuration 2894
Validation 2894
IPv6 Inter-VRF Route Leaking Configuration 2895
R1 Configuration 2895
Validation 2896
OSPF VRF Configuration 2897
Overview 2897
Topology 2897
Configuration IPv4 VRF 2897
R2 2897
Validation 2898
IS-IS VRF Configuration 2899

Overview 2899
Topology 2899
Configuration IPv4 VRF 2899
R1 2899
R2 2900
Validation 2900
IS-IS IPv6 VRF Configuration 2902
Overview 2902
Configuration 2902
Topology 2902
Validation 2903
Maximum Allowed VRF FIB Routes in BGP/MPLS L3VPN Configuration 2905
Overview 2905
Feature Characteristics 2905
Benefits 2905
Configuration 2906
Topology 2906
Running Configuration on PE1 Router 2908
Validation for PE1 2909
Validation for P 2913
Configure PE2 router as follows: 2913
Validation for PE2 2916
Maximum FIB Routes Commands 2920
maximum-fib-routes threshold 2922
maximum-fib-routes 2923
Glossary 2924

Virtual Router Redundancy Protocol Configuration 2925

VRRP Configuration 2926
Terminology 2926
VRRP Process 2926
One Virtual Router 2928
Topology 2929
Configuration 2929
Validation 2930
Two Virtual Routers 2931
Topology 2931
Configuration 2931
Validation 2933
Two Backup Routers 2935
Topology 2935
Configuration 2935
Validation 2938
Disabling the Master/Owner 2939
Interface Tracking 2940
Topology 2941
Configuration 2941

© 2026 IP Infusion Inc. Proprietary

Validation 2944
VRRP-Backward Compatibility 2947
Topology 2948
Configuration 2948
Validation 2949
Redundancy Using VRRP and OSPF: Two Virtual Routers 2950
Topology 2951
Configuration 2952
Verification 2953
VRRP Over MLAG 2954
Topology 2954
Configuration 2954
Validation 2960
Custom VRF Configuration 2961
Topology 2961
Configuration 2961
Validation 2962
Object Tracking Using IP SLA 2966
Overview 2966
Object Tracking Using IP SLA Configuration 2966
Topology 2966
R1 Configuration 2967
R2 Configuration 2969
Validation 2970
VRRP IPv6 Configuration 2973
Topology 2973
Configuration 2973
Owner/Master Router 2973
Backup Router 2974
Host A 2974
Router 1 2975
Validation 2975
Disabling the Master 2976
Validation 2976
GLOBAL IPV6 SUPPORT FOR VRRP 2977
Master Router 2977
Backup Router 2977
Host A 2978
Router 1 2979
Validation 2979

Virtual Router Redundancy Protocol Command Reference 2981

VRRP Commands 2983
accept-mode 2984
Command Syntax 2984
Parameter 2984
Default 2984

Command Mode 2984
Applicability 2984
Examples 2984
advertisement-interval 2985
Command Syntax 2985
Parameter 2985
Default 2985
Command Mode 2985
Applicability 2985
Examples 2985
authentication text 2986
Command Syntax 2986
Parameter 2986
Default 2986
Command Mode 2986
Applicability 2986
Examples 2986
circuit-failover 2987
Command Syntax 2987
Parameters 2987
Default 2987
Command Mode 2987
Applicability 2987
Examples 2987
debug vrrp 2988
Command Syntax 2988
Parameters 2988
Command Mode 2988
Applicability 2988
Examples 2988
disable 2989
Command Syntax 2989
Parameters 2989
Default 2989
Command Mode 2989
Applicability 2989
Example 2989
enable 2990
Command Syntax 2990
Parameters 2990
Default 2990
Command Mode 2990
Applicability 2990
Example 2990
ip pim redundancy 2991
Command Syntax 2991
Parameter 2991

Command Mode 2991
Applicability 2991
Examples 2991
ipv4-exclude-pseudo-header 2992
Command Syntax 2992
Parameter 2992
Default 2992
Command Mode 2992
Applicability 2992
Examples 2992
operational-ip 2993
Command Syntax 2993
Parameters 2993
Default 2993
Command Mode 2993
Applicability 2993
Example 2993
preempt-mode 2994
Command Syntax 2994
Parameters 2994
Default 2994
Command Mode 2994
Applicability 2994
Examples 2994
priority 2995
Command Syntax 2995
Parameter 2995
Default 2995
Command Mode 2995
Applicability 2995
Examples 2995
router vrrp 2996
Command Syntax 2996
Parameters 2996
Default 2996
Command Mode 2996
Applicability 2996
Example 2996
show debugging vrrp 2997
Command Syntax 2997
Parameters 2997
Command Mode 2997
Applicability 2997
Example 2997
show running-config router vrrp 2998
Command Syntax 2998
Parameters 2998

Command Mode 2998
Applicability 2998
Example 2998
show running-config router ipv6 vrrp 2999
Command Syntax 2999
Parameters 2999
Command Mode 2999
Applicability 2999
Example 2999
show vrrp 3000
Command Syntax 3000
Parameters 3000
Command Mode 3000
Applicability 3000
Example 3000
show vrrp (global | ipv4) statistics 3002
Command Syntax 3002
Parameters 3002
Command Mode 3002
Applicability 3002
Example 3002
show vrrp 3003
Command Syntax 3003
Parameters 3003
Command Mode 3003
Applicability 3003
Example 3003
show vrrp summary 3004
Command Syntax 3004
Parameters 3004
Command Mode 3004
Applicability 3004
Example 3004
snmp restart vrrp 3005
Command Syntax 3005
Parameters 3005
Default 3005
Command Mode 3005
Applicability 3005
Examples 3005
switch-back-delay 3006
Command Syntax 3006
Parameters 3006
Command Mode 3006
Default 3006
Applicability 3006
Examples 3006

track decrement 3007
Command Syntax 3007
Parameters 3007
Command Mode 3007
Applicability 3007
Example 3007
undebug vrrp 3008
Command Syntax 3008
Parameters 3008
Command Mode 3008
Applicability 3008
Examples 3008
virtual-ip 3009
Command Syntax 3009
Parameters 3009
Default 3009
Command Mode 3009
Applicability 3009
Example 3009
vrrp compatible-v2 3010
Command Syntax 3010
Parameters 3010
Default 3010
Command Mode 3010
Applicability 3010
Examples 3010
vrrp ipv4-exclude-pseudo-header 3011
Command Syntax 3011
Parameters 3011
Default 3011
Command Mode 3011
Applicability 3011
Examples 3011
vrrp vmac 3012
Command Syntax 3012
Parameters 3012
Default 3012
Command Mode 3012
Applicability 3012
Example 3012
VRRP v6 Commands 3013
advertisement-interval 3014
Command Syntax 3014
Parameter 3014
Command Mode 3014
Applicability 3014
Examples 3014

circuit-failover 3015
Command Syntax 3015
Parameters 3015
Command Mode 3015
Applicability 3015
Examples 3015
disable 3016
Command Syntax 3016
Parameters 3016
Command Mode 3016
Applicability 3016
Example 3016
enable 3017
Command Syntax 3017
Parameters 3017
Command Mode 3017
Applicability 3017
Example 3017
preempt-mode 3018
Command Syntax 3018
Parameters 3018
Default 3018
Command Mode 3018
Applicability 3018
Example 3018
priority 3019
Command Syntax 3019
Parameter 3019
Default 3019
Command Mode 3019
Applicability 3019
Examples 3019
router ipv6 vrrp 3020
Command Syntax 3020
Parameters 3020
Command Mode 3020
Applicability 3020
Example 3020
router ipv6 vrrp vlan 3021
Command Syntax 3021
Parameters 3021
Command Mode 3021
Applicability 3021
Examples 3021
show running-config vrrpv6 3022
Command Syntax 3022
Parameters 3022

Command Mode 3022
Applicability 3022
Example 3022
virtual-ipv6 3023
Command Syntax 3023
Parameters 3023
Command Mode 3023
Applicability 3023
Examples 3023

Unidirectional Link Detection Configuration 3024

Unidirectional Link Detection Configuration 3025
Topology 3025
S1 3025
S2 3026
Validation 3026
Enable UDLD under bridge-group 3027
S1 3027
S2 3027
Validation 3028

Unidirectional Link Detection Command Reference 3029

Unidirectional Link Detection Commands 3030
udld 3031
Command Syntax 3031
Parameters 3031
Default 3031
Command Mode 3031
Applicability 3031
Examples 3031
udld message-time 3032
Command Syntax 3032
Parameters 3032
Default 3032
Command Mode 3032
Applicability 3032
Examples 3032
udld mode 3033
Command Syntax 3033
Parameters 3033
Default 3033
Command Mode 3033
Applicability 3033
Examples 3033
udld state 3034
Command Syntax 3034
Parameters 3034
Default 3034

© 2026 IP Infusion Inc. Proprietary

Command Mode 3034
Applicability 3034
Examples 3034
show udld 3035
Command Syntax 3035
Parameters 3035
Command Mode 3035
Applicability 3035
Examples 3035
show udld interface 3036
Command Syntax 3036
Parameters 3036
Command Mode 3036
Applicability 3036
Examples 3036

Bridged Virtual Interface 3037

QoS and ACL over Bridge Virtual Interface (BVI) 3038
Overview 3038
Limitations 3038
Prerequisites 3039
Configuration 3039
Topology 3039
Configure QoS and ACL on BVI Interface in R1 Router 3039
Configure QoS and ACL on BVI Interface in R2 Router 3040
Running Configuration 3041
QoS and ACL on BVI Interface in R1 3041
QoS and ACL on BVI Interface in R2 3041
Validation 3042
R1 3042
R2 3043
Glossary 3044

© 2026 IP Infusion Inc. Proprietary

PREFACE

About this Guide

This guide describes how to configure Layer 3 in OcNOS.

Audience

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

Conventions

The Table 1 table 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

IP Infusion Product Release Version

Each integer in release numbers indicates Major, Minor, and Maintenance release versions. Build numbers that follow the release numbers are for internal tracking and verification of the software build process and are visible to customers as part of the software version number.

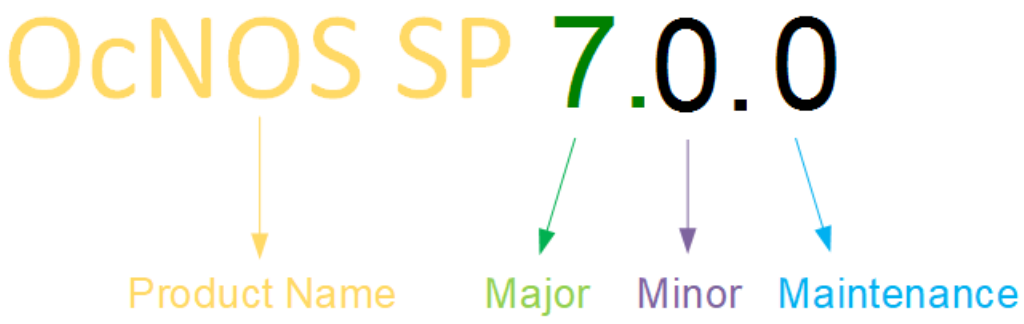

Product Name: IP Infusion Product Family

Major Version: New customer-facing functionality that represents a significant change to the code base; including a significant marketing change or direction in the product.

Minor Version: Enhancements or extensions to existing features, changes to address external needs, or internal improvements to satisfy new sales regions or marketing initiatives.

Maintenance Version: A collection of product bugs or issues usually scheduled every 30 or 60 days, based on the number of issues.

Related Documentation

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

Feature Availability

Each OcNOS SKU contains a set of supported features. For a list of available features based on the SKU that you purchased, refer to the *Feature Matrix* .

Migration Guide

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

IP Maestro Support

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

Technical Support

IP Infusion maintains an online technical support site that provides a variety of technical support programs for licensed OcNOS customers at the Technical Assistance Center.

Customers and partners enjoy full access to the support website. The site allows customers and partners to open technical support calls, update open calls with new information, and review the status of open or closed calls. The password-protected site includes technical documentation, Release Notes, and descriptions of service offerings.

Technical Sales

Contact the IP Infusion sales representative for more information about the OcNOS solution.

Technical Documentation

For core commands and configuration procedures, visit: Product Documentation.

For training videos, visit: OcNOS Free Training Videos.

For a list of supported platforms and SKUs of OcNOS features, refer to the OcNOS Feature Matrix.

© 2026 IP Infusion Inc. Proprietary

Documentation Disclaimer

The global documentation site is evolving to provide an enhanced website user experience for select topics included in this release. Some guides are now available outside the existing documentation library and can be accessed directly from custom documentation landing pages. These guides offer robust in-built search functionality.

For the latest documentation, visit the product-specific documentation landing page and select the relevant guide.

Comments

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

© 2026 IP Infusion Inc. Proprietary

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.

Chapter Organization

The chapters in command references are organized as described in Command Description Format (page 173).

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

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 authentication-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

© 2026 IP Infusion Inc. Proprietary

Syntax Conventions

Table 2 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 (page 172)	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>}
[]	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 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

The Table 4 table 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 (page 179)
Applicability	The command introduced in a specific release version and modified or updated in subsequent versions.
Example	An example of the command being executed

© 2026 IP Infusion Inc. Proprietary

Keyboard Operations

The Table 5 table 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
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

© 2026 IP Infusion Inc. Proprietary

Show Command Modifiers

Note: The show command output included in the guides is for illustration purposes only. Based on the combination of features enabled and ongoing enhancements made to the commands, the output for these commands may vary. For instance, the actual command output may differ depending on the software version, configuration, and platform. Field names, values, and formats are subject to change.

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
!
```

© 2026 IP Infusion Inc. Proprietary

```
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 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.

© 2026 IP Infusion Inc. Proprietary

Command Modes

Commands are grouped into modes arranged in a hierarchy. Each mode has its own set of commands. The table below lists the command modes common to all protocols.

Table 7. Common Command Modes

Name	Description
`Execution 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 execution 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.

© 2026 IP Infusion Inc. Proprietary

Command Mode Tree

The diagram below shows the common command mode hierarchy.

Figure 1. Common command modes

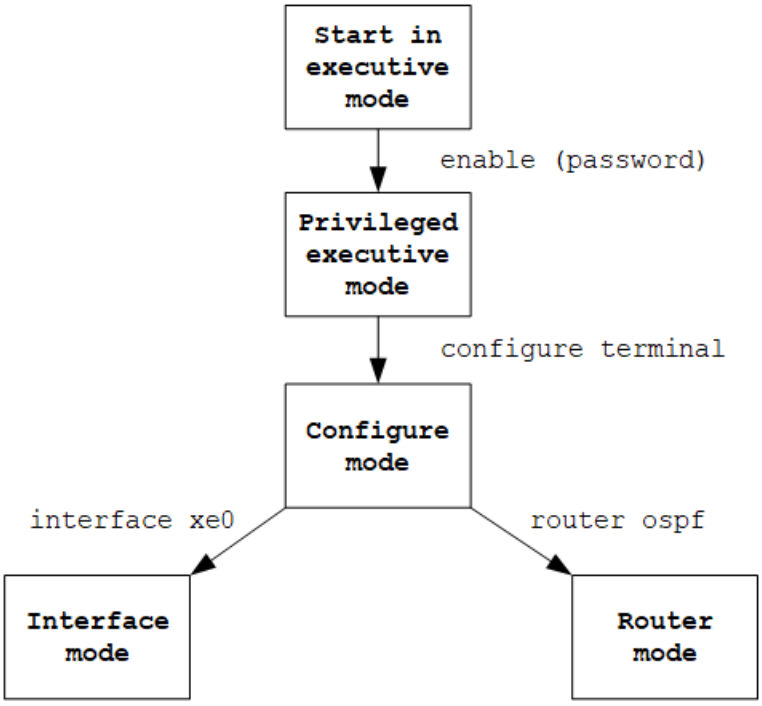

To change modes:

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.

© 2026 IP Infusion Inc. Proprietary

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 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 . 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.

© 2026 IP Infusion Inc. Proprietary

LAYER 3 UNICAST CONFIGURATION

Forwarding Plane Load Balancing 184
Enable Load Balancing 184
Configuration 184
Validation 186
VLAN Interfaces 187
Overview 187
Topology 187
Create a VLAN Interface 188
Layer 3 Link Aggregation 191
Overview 191
Configuration 191
IPv4Static Routes 196
Configuration 196
IPv6 Static Routs 201
Configuration 201
Static Route Discard Configuration 206
Overview 206
IPv4 Route Discard 206
Validation 208
IPv6 Route Discard 209
Validation 212
Static Route Object Tracking using IP SLA 214
Overview 214
Topology 214
Static Route Object Tracking IPv4 Configuration 214
Static Route Object Tracking IPv6 Configuration 220
Route Monitor 224
Overview 224
Configuration 224
Implementation Examples 232
Troubleshooting 233
Hybrid Switch Router Configuration 235
Overview 235
Routing and Switching 235
System Configuration 236
Configuring Layer 2 Interfaces 238
Configuring Layer 3 Interfaces 238

© 2026 IP Infusion Inc. Proprietary

Policy Based Routing Configuration 240
Topology 240
Configurations - PBR IPv4 241
Configuration - PBR IPv6 245
PBR with Nexthop Tracking 250
IPv4 Configurations for PBR with Nexthop Tracking 250
IPv6 Configurations for PBR with Nexthop Tracking 252
TOS based Queue Distribution Configuration 255
Precedence-to-queue 255
Precedence-to-precedence 256
Internet Protocol SLA Configuration 258
Topology 258
Configuration 258
Validation 259
EVPN AF Route-Maps and Route Filtering 261
Overview 261
Feature Characteristics 261
Benefits 261
Prerequisites 261
Configuration 262
Implementation Examples 271
Commands 271
Glossary 275

© 2026 IP Infusion Inc. Proprietary

Forwarding Plane Load Balancing

OcNOS uses Forwarding Plane Load Balancing when the kernel supports Equal Cost Multipath (ECMP). OcNOS installs the maximum number of ECMP routes supported by the kernel. This allows for load balancing to be performed with more than one nexthop to reach a destination. If the router receives and installs multiple paths with the same administrative distance and cost to a destination, load-balancing is possible.

Ideally, multiple nexthops have different interfaces to the destination, but this is not mandatory. The algorithm for distributing traffic across ECMP routes is dependent on the kernel, and typically based on the protocol, source address, destination address, and port.

Enable Load Balancing

The following example illustrates how to enable Equal Cost Multipath (ECMP), and configure a routing protocol (OSPF is used in this example) for load balancing. However, this example will not work if the kernel does not support load balancing. In this topology, R1, R2, and R3 are three Linux routers connected to each other. R1 can reach R3 through two links available to R2.

Topology

Figure 2. Load Balancing Topology

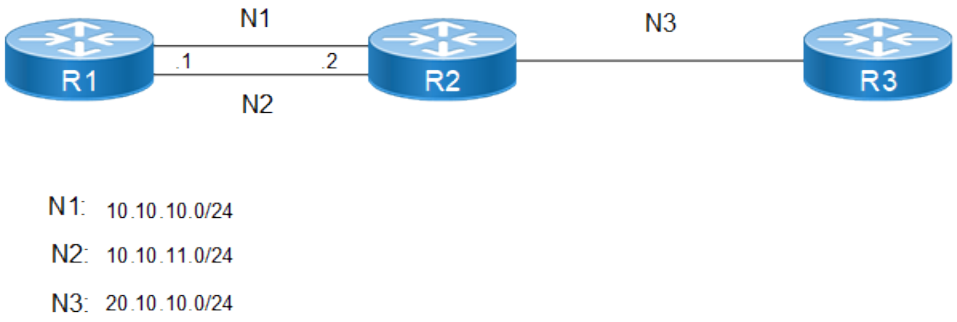

Configuration

R1 - NSM

- Enable multipath support and set the maximum number of paths to be installed in the Forwarding Information Base (FIB):

`#configure terminal`	Enter configure mode
`(config)#maximum-paths 2`	Create maximum paths 2
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit interface mode

© 2026 IP Infusion Inc. Proprietary

Note: User can save and reboot to bring changes into effect.

R1 - OSPF

- Configure OSPF on all interfaces on R1, R2, and R3.
  R1 learns about R3 through 2 nexthops (both networks N1 and N2).

`#configure terminal`	Enter configure mode
`(config)#router ospf 1`	Enter ospf router mode
`(config-router)#router 1.1.1.1`	Specify router id
`(config-router)#net 10.11.1.1/30 a 0`	Configure ospf network in area 0
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#end`	Exit interface mode

R2 - OSPF

`#configure terminal`	Enter configure mode
`(config)#router ospf 1`	Enter ospf router mode
`(config-router)#router 2.2.2.2`	Specify router id
`(config-router)#net 10.11.1.2/30 a 0`	Configure ospf network in area 0
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#end`	Exit interface mode

R3 - OSPF

`#configure terminal`	Enter configure mode
`(config)#router ospf 1`	Enter ospf router mode
`(config-router)#router 3.3.3.3`	Specify router id
`(config-router)#net 10.1.1.2/30 a 0`	Configure ospf network in area 0
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#end`	Exit interface mode

© 2026 IP Infusion Inc. Proprietary

Validation

R1 - OSPF

Run the `show ip ospf route` command on R1. The OSPF routing table displays that it can reach R3 through both of the nexthops:

```
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  10.10.10.0/24 [10] is directly connected, eth1, Area 0.0.0.0
C  10.10.11.0/24 [10] is directly connected, eth2, Area 0.0.0.0
O  20.10.10.0/24 [11] via 10.10.10.3, eth1, Area 0.0.0.0
via 10.10.11.3, eth2, Area 0.0.0.0
```

Run the `show ip route` command on R1. It displays that R1 has installed both nexthops to reach R3 in the NSM routing table:

```
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, 00:17:08
C     10.10.11.0/24 is directly connected, eth2, 00:16:49
O     20.10.10.0/24 [110/11] via 10.10.11.3, eth2, 00:08:52
             [110/11] via 10.10.10.3, eth1
C     127.0.0.0/8 is directly connected, lo, 00:25:21
C     192.168.52.0/24 is directly connected, eth0, 00:25:16

Gateway of last resort is not set
```

© 2026 IP Infusion Inc. Proprietary

VLAN Interfaces

This section contains examples for configuring VLAN interfaces.

Overview

Several Virtual LAN (VLAN) interfaces can be configured on a single Ethernet interface. Once created, a VLAN interface functions the same as any physical interface.

NSM recognizes VLAN interfaces as physical interfaces. Once VLAN interfaces are created in the kernel, and IP addresses are assigned to them, OcNOS commands can be used to configure and display VLAN interfaces the same as any physical interface. OcNOS routing protocols, such as OSPF and BGP can run across networks using VLAN interfaces.

Two systems with physical connectivity (either directly connected or connected through a switch), can communicate with each other via VLAN interfaces that have the same VLAN IDs and belong to the same network.

If the physical interfaces are not directly connected to a switch, the corresponding ports on the switch must be configured as trunks, and should not be associated to any VLANs in the switch. The commands to configure switch ports as trunks depend on the type of the switch, and are beyond the scope of this document.

Topology

Figure 3 is used to describe VLAN interface configuration. In this example, there are two routers, R1 and R2, and the eth1 interface of R1 is connected directly to eth2 using an ethernet cable.

Figure 3. VLAN Connections

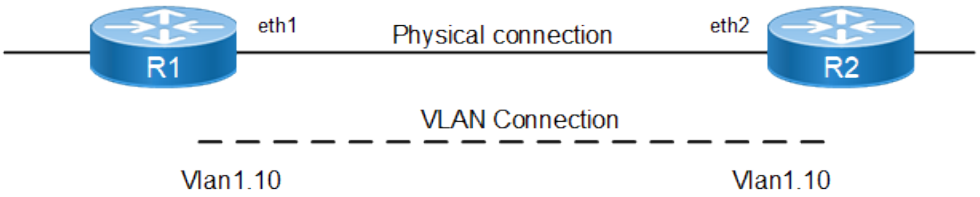

The `vlan1.10` VLAN interface is created on R1, and `vlan1.10` is created on R2. The VLAN interfaces are configured in the same network: R1 and R2 can reach each other using the VLAN connection.

Note: A VLAN ID of both VLAN interfaces is the same (10). Two systems with different VLAN IDs cannot communicate, even if they are in the same network, since a VLAN ID tags packets sent on a VLAN interface.

© 2026 IP Infusion Inc. Proprietary

Create a VLAN Interface

When a VLAN interface is configured, a Layer 3 interface based on the bridge-group number and VLAN ID is created. This Layer 3 interface is advertised to all the Layer 3 protocols.

`#configure terminal`	Enter configure mode.
`(config)#bridge 1 protocol mstp`	Create a MSTP bridge.
`(config)#vlan database`	Enter VLAN mode.
`(config-vlan)#vlan 10 bridge 1`	Enable VLAN 10 on bridge 1.
`(config-vlan)#exit`	Exit VLAN mode.
`(config)#interface eth1`	Enter interface mode
`(config-if)#switchport`	Configure interface as Layer2 interface.
`(config-if)#bridge-group 1`	Associate bridge group 1.
`(config-if)#switchport mode trunk`	Configure interface eth1 as Layer2 trunk mode.
`(config-if)#switchport trunk allowed vlan add 10`	Associate VLAN 10 as trunk port.
`(config-if)#commit`	Commit the candidate configuration to the running
`(config-if)#exit`	Exit interface mode.

Add IP Addresses to VLAN Interface

In NSM, you can add or remove IP addresses from VLAN interfaces, like normal interfaces. Using IMISH type:

`#configure terminal`	Enter configure mode
`(config)#interface vlan1.10`	Enter interface mode
`(config-if)#ip address 192.168.1.50/24`	Assign ip address to vlan interface
`(config-if)#commit`	Commit the candidate configuration to the running
`(config-if)#end`	Exit interface mode.

Display VLAN Interfaces

In OcNOS, VLAN interfaces appear as any physical interfaces, in the `show running-config` or the `show ip interface brief` outputs, and can be configured as any other interface.

The following is a sample output of the `show ip interface brief` command on R1.

Note: The IP address of interface vlan1.10 has correctly been changed by NSM:

```
#show ip interface brief

'*' - address is assigned by dhcp client
```

```
Interface            IP-Address      Admin-Status          Link-Status
eth0                 10.12.56.26     up                    up
lo                   127.0.0.1       up                    up
lo.management        127.0.0.1       up                    up
vlan1.1              unassigned      up                    up
vlan1.10             192.168.1.50    up                    up
xe1                  unassigned      up                    up
xe2                  unassigned      up                    up
xe3                  unassigned      up                    down
xe4                  unassigned      up                    down
xe5                  unassigned      up                    down
xe6                  unassigned      up                    down
xe7                  unassigned      up                    down
xe8                  unassigned      up                    down
xe9                  unassigned      up                    down
xe10                 unassigned      up                    down
xe11                 unassigned      up                    down
xe12                 unassigned      up                    down
xe13                 unassigned      up                    up
xe14                 unassigned      up                    down
xe15                 unassigned      up                    down
xe16                 unassigned      up                    up
xe17                 unassigned      up                    down
xe18                 unassigned      up                    down
xe19                 unassigned      up                    down
xe20                 unassigned      up                    down
xe21                 unassigned      up                    down
xe22                 unassigned      up                    down
xe23                 unassigned      up                    down
xe24                 unassigned      up                    down
xe25                 unassigned      up                    down
xe26                 unassigned      up                    down
xe28                 unassigned      up                    down
xe29                 unassigned      up                    down
xe30                 unassigned      up                    down
xe31                 unassigned      up                    down
xe32                 unassigned      up                    up
xe33                 unassigned      up                    down
xe34                 unassigned      up                    down
xe35                 unassigned      up                    down
xe36                 unassigned      up                    down
xe37                 unassigned      up                    up
xe38                 unassigned      up                    down
xe39                 unassigned      up                    down
xe40                 unassigned      up                    down
xe41                 unassigned      up                    up
xe42                 unassigned      up                    down
xe43                 unassigned      up                    down
xe44                 unassigned      up                    down
xe45                 unassigned      up                    up
xe46                 unassigned      up                    down
xe47                 unassigned      up                    down
xe48                 unassigned      up                    up
xe49/1               unassigned      up                    up
xe49/2               unassigned      up                    down
xe49/3               unassigned      up                    down
xe49/4               unassigned      up                    down
xe50/1               unassigned      up                    up
xe50/2               unassigned      up                    down
xe50/3               unassigned      up                    down
xe50/4               unassigned      up                    down
xe51/1               unassigned      up                    down
xe51/2               unassigned      up                    down
xe51/3               unassigned      up                    down
xe51/4               unassigned      up                    down
xe52/1               unassigned      up                    down
xe52/2               unassigned      up                    down
```

© 2026 IP Infusion Inc. Proprietary

```
xe52/3              unassigned      up                  down
xe52/4              unassigned      up                  down
xe53/1              unassigned      up                  up
xe53/2              unassigned      up                  down
xe53/3              unassigned      up                  down
xe53/4              unassigned      up                  down
xe54/1              unassigned      up                  up
xe54/2              unassigned      up                  down
xe54/3              unassigned      up                  down
xe54/4              unassigned      up                  down
```

Below is the NSM routing table, which shows the connected network `192.168.1.0/24` of `vlan1.10`. These interfaces will now act as any physical interfaces, and all routing protocols will run across this 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+F27 area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C     127.0.0.0/8 is directly connected, lo, 00:16:43
C     192.168.1.0/24 is directly connected, vlan1.10, 00:02:05
C     192.168.52.0/24 is directly connected, eth0, 00:16:39

Gateway of last resort is not set
```

Layer 3 Link Aggregation

Overview

This section contains a complete sample Link Aggregation Control Protocol (LACP) configuration (L3 LAG).

Link Aggregation is the method of combining individual physical network interfaces or ports to increase the capacity of the link to support and sustain beyond the individual port capability. Features like Spanning Tree, VLAN, FDB, Multicast operate on both physical ports as well as Link Aggregated Logical Ports. It bundles all of the controller's distribution system ports into a single 802.3ad port channel, thereby reducing the number of IP addresses needed to configure the ports on your controller. When LAG is enabled, the system dynamically manages port redundancy and load balances access points transparently to the user.

LACP is based on the 802.3ad IEEE specification. It allows bundling of several physical interfaces to form a single logical channel providing enhanced performance and redundancy. The aggregated interface is viewed as a single link to each switch. The spanning tree views it as one interface and not as two or three interfaces. When there is a failure in one physical interface, the other interfaces stay up and there is no disruption.

The OcNOS LACP implementation supports the aggregation of a maximum of six physical Ethernet links into a single logical channel.

Topology

In this example, 3 links are configured between the two switches R1 and R2. These three links are assigned the same administrative key (1) so that they aggregate to form a single channel 1. They are viewed by STP as one interface.

Figure 4. L3_LAG Topology

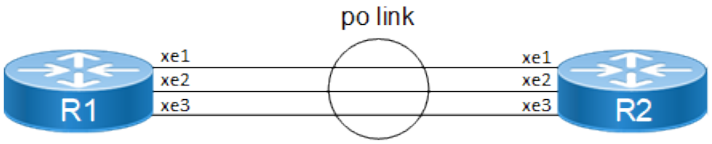

Configuration

R1

`R1#configure terminal`	Enter configure mode.
`R1(config)#interface po10`	Enter interface mode.
`R1(config-if)#ip address 1.1.1.1/24`	Assigning IP Address to PO Interface
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#lacp system-priority 20000`	Set the system priority of this switch. This priority is used for determining the system that is responsible

© 2026 IP Infusion Inc. Proprietary

	for resolving conflicts in the choice of aggregation groups. A lower numerical value has a higher priority.
`R1(config)#interface xe1`	Enter interface mode.
`R1(config-if)#no switchport`	Making Interface as L3 Port (This command will remove if switchport configuration is present).
`R1(config-if)#channel-group 10 mode active`	Add this interface to channel group 10 and enable link aggregation so that it can be selected for aggregation by the local system.
`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.
`R1(config-if)#no switchport`	Making Interface as L3 Port (This command will remove if switchport configuration is present).
`R1(config-if)#channel-group 10 mode active`	Add this interface to channel group 10 and enable link aggregation so that it can be selected for aggregation by the local system.
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#interface xe3`	Enter interface mode.
`R1(config-if)#no switchport`	Making Interface as L3 Port (This command will remove if switchport configuration is present).
`R1(config-if)#channel-group 10 mode active`	Add this interface to channel group 10 and enable link aggregation so that it can be selected for aggregation by the local system.
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration
`R1(config-if)#exit`	Exit interface mode.

R2

`R2#configure terminal`	Enter configure mode.
`R2(config)#interface po10`	Enter interface mode.
`R2(config-if)#ip address 1.1.1.2/24`	Assigning IP Address to PO Interface
`R2(config-if)#commit`	Commit the candidate configuration to the running configuration
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#lacp system-priority 20000`	Set the system priority of this switch. This priority is used for determining the system that is responsible for resolving conflicts in the choice of aggregation groups. A lower numerical value has a higher priority.

© 2026 IP Infusion Inc. Proprietary

`R2(config)#interface xe1`	Enter interface mode.
`R2(config-if)#no switchport`	Making Interface as L3 Port (This command will remove if switchport configuration is present).
`R2(config-if)#channel-group 10 mode active`	Add this interface to channel group 10 and enable link aggregation so that it can be selected for aggregation by the local system.
`R1(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.
`R2(config-if)#no switchport`	Making Interface as L3 Port (This command will remove if switchport configuration is present).
`R2(config-if)#channel-group 10 mode active`	Add this interface to channel group 10 and enable link aggregation so that it can be selected for aggregation by the local system.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#interface xe3`	Enter interface mode.
`R2(config-if)#no switchport`	Making Interface as L3 Port (This command will remove if switchport configuration is present).
`R2(config-if)#channel-group 10 mode active`	Add this interface to channel group 10 and enable link aggregation so that it can be selected for aggregation by the local system.
`R2(config-if)#commit`	Commit the candidate configuration to the running configuration
`R2(config-if)#exit`	Exit interface mode.

Static Channel-group

R1

`R1#configure terminal`	Enter configure mode
`R1(config)#interface sa12`	Enter interface mode
`R1(config-if)#no switchport`	Making Interface as L3 Port (This command will remove if switchport configuration is present).
`R1(config-if)#ip address 2.2.2.1/24`	Assigning IP Address to PO 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
`R1(config-if)# static-channel-group 12`	Add this interface to channel group 12 and enable link aggregation so that it can be selected for aggregation by the local system.

© 2026 IP Infusion Inc. Proprietary

R1(config-if)#commit	Commit the candidate configuration to the running configuration
R1(config-if)#exit	Exit interface mode

R2

R2#configure terminal	Enter configure mode
R2(config)#interface sa12	Enter interface mode
R2(config-if)#no switchport	Making Interface as L3 Port (This command will remove if switchport configuration is present).
R2(config-if)#ip address 2.2.2.2/24	Assigning IP Address to PO Interface
R1(config-if)#commit	Commit the candidate configuration to the running configuration
R2(config-if)#exit	Exit interface mode
R2(config)#interface xe1	Enter interface mode
R2(config-if)# static-channel-group 12	Add this interface to channel group 12 and enable link aggregation so that it can be selected for aggregation by the local system.
R1(config-if)#commit	Commit the candidate configuration to the running configuration
R2(config-if)#exit	Exit interface mode

Validation

Show etherchannel detail, show etherchannel summary.

```
#show etherchannel summary
% Aggregator po10 100010
% Aggregator Type: Layer3
% Admin Key: 0010 - Oper Key 0010
%   Link: xe1 (10049) sync: 1
%   Link: xe2 (10050) sync: 1
%   Link: xe3 (10051) sync: 1

#sh etherchannel detail
% Aggregator po10 100010
% Aggregator Type: Layer3
%  Mac address: 14:18:77:5d:5c:01
%  Admin Key: 0010 - Oper Key 0010
%  Actor LAG ID- 0x4e20,14-18-77-01-5c-00,0x000a
%  Receive link count: 3 - Transmit link count: 3
%  Individual: 0 - Ready: 1
%  Partner LAG ID- 0x4e20,14-18-77-01-73-00,0x000a
%   Link: xe1 (10049) sync: 1
%   Link: xe2 (10050) sync: 1
%   Link: xe3 (10051) sync: 1
% Collector max delay: 5

#sh etherchannel 10
% Aggregator po10 100010 Admin Key: 0010 - Oper Key 0010
% Partner LAG ID: 0x4e20,14-18-77-01-73-00,0x000a
% Partner Oper Key 0010

#sh etherchannel
```

```
% Lacp Aggregator: po10
% Member:
   xe1
   xe2
   xe3
#show static-channel-group
%Static Aggregator: sa12
% Member Status
%   xe1    up
%   xe2    up
%   xe3    up
```

IPv4Static Routes

This section 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.

For details about the commands used in these examples, see the Unicast Routing Information Base Command Reference.

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 5. Basic Static Route

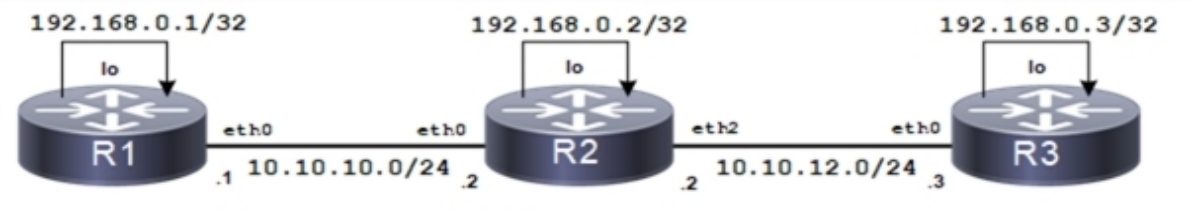

In all three routes, interface `eth0` of router R2 is the gateway. Router R3 is configured with a default static route that is equivalent to configuring separate static routes with the same gateway or nexthop address. Router R2 has two routes, one for each of the remote routers' loopback address.

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)#commit`	Commit the candidate configuration to the running configuration
`(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

© 2026 IP Infusion Inc. Proprietary

	same static route for individual addresses. See the configuration of R3.
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit configure mode

R2

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Enter interface mode.
`(config-if)#ip address 192.168.0.2/32 secondary`	Configure the IP address on this interface, and specify a 32-bit mask, making it a host address.
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(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.
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit configure mode

R3

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Enter interface mode.
`(config-if)#ip address 192.168.0.3/32 secondary`	Configure the IP address on this interface, and specify a 32-bit mask, making it a host address.
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(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.
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit configure mode

Validation

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

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
       * - 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

```
#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       10.10.12.0/24 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

#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      : 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

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

#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
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

```
#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       10.10.12.0/24 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

#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      : 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
```

IPv6 Static Routs

This example shows complete configuration to enable IPv6 static routing in a simple network topology.

Note: IPv6 static route with interface alone as gateway (without gateway IPv6 address) is not supported.

Topology

Figure 6. IPv6 static routing

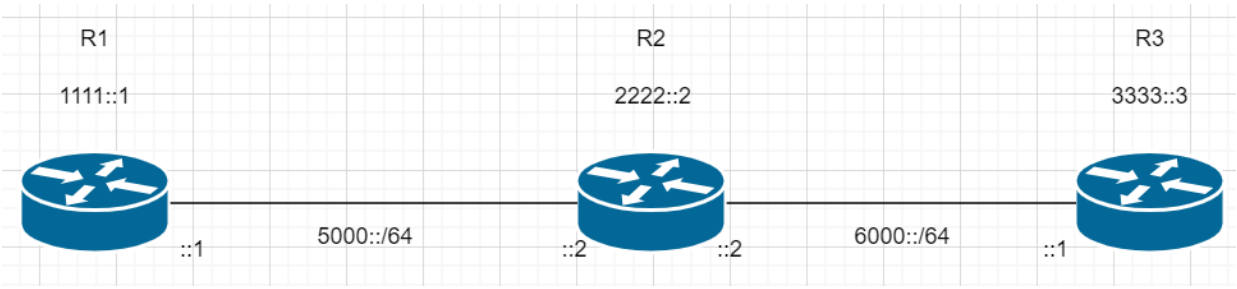

Configuration

R1

`R1#conf t`	Enter Configure mode.
`R1(config)#interface lo`	Enter interface mode.
`R1(config-if)#ipv6 address 1111::1/128`	Configure IPv6 address
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#ipv6 route 6000::/64 5000::2`	Configure IPv6 Static route to reach remote network with R2 as nexthop
`R1(config)#ipv6 route 2222::2/128 5000::2`	Configure IPv6 static route to reach R2 lo with R2 as nexthop
`R1(config)#ipv6 route 3333::3/128 5000::2`	Configure IPv6 static route to reach R3 lo with R2 as nexthop
`R1(config)#commit`	Commit the candidate configuration to the running configuration
`R1(config)#exit`	Exit configure mode

© 2026 IP Infusion Inc. Proprietary

R2

`R2#conf t`	Enter Configure mode.
`R2(config)#int lo`	Enter interface mode.
`R2(config-if)#ipv6 address 2222::2/128`	Configure IPv6 address
`R2(config)#commit`	Commit the candidate configuration to the running configuration
`R2(config)#exit`	Exit configure mode
`R2(config)#ipv6 route 1111::1/128 5000::1`	Configure IPv6 static route to reach R1 lo with R1 as nexthop
`R2(config)#ipv6 route 3333::3/128 6000::1`	Configure IPv6 static route to reach R3 lo with R3 as nexthop
`R2(config)#commit`	Commit the candidate configuration to the running configuration
`R2(config)#exit`	Exit configure mode

R3

`R3#conf t`	Enter Configure mode.
`R3(config)#int lo`	Enter interface mode.
`R3(config-if)#ipv6 add 3333::3/128`	Configure IPv6 address
`R3(config)#commit`	Commit the candidate configuration to the running configuration
`R3(config-if)#exit`	Exit interface mode.
`R3(config)#ipv6 route ::/0 6000::2`	Configure Default IPv6 Static route with R2 as nexthop
`R3(config)#commit`	Commit the candidate configuration to the running configuration
`R3(config)#exit`	Exit configure mode

Validation

Show ipv6 route, show ipv6 route summary, show ipv6 route database

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, 20:51:02
C      1111::1/128 via ::, lo, 00:01:13
```

```
S      2222::2/128 [1/0] via 5000::2, xe3, 00:00:32
S      3333::3/128 [1/0] via 5000::2, xe3, 00:00:13
C      5000::/64 via ::, xe3, 00:01:42
S      6000::/64 [1/0] via 5000::2, xe3, 00:00:54
C      fe80::/64 via ::, ce45, 01:45:19
R1#show ipv6 route summary
IPv6 routing table name is Default-IPv6-Routing-Table(0)
IPv6 routing table maximum-paths : 8
Total number of IPv6 routes      : 7
Total number of IPv6 paths       : 7
Pending routes (due to route max reached): 0
Route Source    Networks
connected       4
static          3
Total           7
FIB             7

ECMP statistics (active in ASIC):
---------------------------------
 Total number of IPv6 ECMP routes   : 0
 Total number of IPv6 ECMP paths    : 0
R1#
R1#show ipv6 route database
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
       > - selected route, * - FIB route,p - stale info
Timers: Uptime

IP Route Table for VRF "default"
C   *> ::1/128 via ::, lo, 20:51:19
C   *> 1111::1/128 via ::, lo, 00:01:30
S   *> 2222::2/128 [1/0] via 5000::2, xe3, 00:00:49
S   *> 3333::3/128 [1/0] via 5000::2, xe3, 00:00:30
C   *> 5000::/64 via ::, xe3, 00:01:59
S   *> 6000::/64 [1/0] via 5000::2, xe3, 00:01:11
C   *> fe80::/64 via ::, ce45, 01:45:36
C      fe80::/64 via ::, ce44, 01:45:36
C      fe80::/64 via ::, xe39, 01:45:36
C      fe80::/64 via ::, xe32, 01:45:36
C      fe80::/64 via ::, xe29, 01:45:36
C      fe80::/64 via ::, xe13, 01:45:36
C      fe80::/64 via ::, ce46, 03:56:36
C      fe80::/64 via ::, ce43, 03:56:36
C      fe80::/64 via ::, xe25, 03:56:36
C      fe80::/64 via ::, xe23, 03:56:36
C      fe80::/64 via ::, xe3, 03:56:36
C      fe80::/64 via ::, xe34, 20:41:33
C      fe80::/64 via ::, xe33, 20:41:33
C      fe80::/64 via ::, xe36, 20:50:48
C      fe80::/64 via ::, xe22, 20:50:48
C      fe80::/64 via ::, xe21, 20:50:48
C      fe80::/64 via ::, xe10, 20:50:48
C      fe80::/64 via ::, xe9, 20:50:48
```

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, 03:59:56
S      1111::1/128 [1/0] via 5000::1, xe3, 00:00:46
C      2222::2/128 via ::, lo, 00:01:27
S      3333::3/128 [1/0] via 6000::1, xe5, 00:00:26
C      5000::/64 via ::, xe3, 00:01:52
C      6000::/64 via ::, xe5, 00:01:10
C      fe80::/64 via ::, vlan1.2, 01:17:00
R2#
R2#show ipv6 route summary
IPv6 routing table name is Default-IPv6-Routing-Table(0)
IPv6 routing table maximum-paths : 8
Total number of IPv6 routes      : 7
Total number of IPv6 paths       : 7
Pending routes (due to route max reached): 0
Route Source    Networks
connected       5
static          2
Total           7
FIB             7

ECMP statistics (active in ASIC):
---------------------------------
 Total number of IPv6 ECMP routes   : 0
 Total number of IPv6 ECMP paths    : 0
R2#show ipv6 route database
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
       > - selected route, * - FIB route,p - stale info
Timers: Uptime

IP Route Table for VRF "default"
C   *> ::1/128 via ::, lo, 04:00:02
S   *> 1111::1/128 [1/0] via 5000::1, xe3, 00:00:52
C   *> 2222::2/128 via ::, lo, 00:01:33
S   *> 3333::3/128 [1/0] via 6000::1, xe5, 00:00:32
C   *> 5000::/64 via ::, xe3, 00:01:58
C   *> 6000::/64 via ::, xe5, 00:01:16
C   *> fe80::/64 via ::, vlan1.2, 01:17:06
C      fe80::/64 via ::, vlan1.1, 01:17:06
C      fe80::/64 via ::, xe29, 01:48:22
C      fe80::/64 via ::, xe27, 01:48:22
C      fe80::/64 via ::, ce47, 03:59:22
C      fe80::/64 via ::, ce46, 03:59:22
C      fe80::/64 via ::, ce45, 03:59:22
C      fe80::/64 via ::, ce43, 03:59:22
C      fe80::/64 via ::, xe42, 03:59:22
C      fe80::/64 via ::, xe41, 03:59:22
C      fe80::/64 via ::, xe34, 03:59:22
C      fe80::/64 via ::, xe33, 03:59:22
C      fe80::/64 via ::, xe32, 03:59:22
C      fe80::/64 via ::, xe31, 03:59:22
C      fe80::/64 via ::, xe25, 03:59:22
C      fe80::/64 via ::, xe23, 03:59:22
C      fe80::/64 via ::, xe5, 03:59:22
C      fe80::/64 via ::, xe3, 03:59:22
R2#
```

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"
S      ::/0 [1/0] via 6000::2, xe5, 00:00:07
C      ::1/128 via ::, lo, 20:46:35
C      3333::3/128 via ::, lo, 00:00:57
C      6000::/64 via ::, xe5, 00:00:46
C      fe80::/64 via ::, ce43, 01:50:07
R3#show ipv6 route summary
IPv6 routing table name is Default-IPv6-Routing-Table(0)
IPv6 routing table maximum-paths : 8
Total number of IPv6 routes      : 5
Total number of IPv6 paths       : 5
Pending routes (due to route max reached): 0
Route Source    Networks
connected       4
static          1
Total           5
FIB             5

ECMP statistics (active in ASIC):
---------------------------------
 Total number of IPv6 ECMP routes   : 0
 Total number of IPv6 ECMP paths    : 0
R3#
R3#show ipv6 route database
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
       > - selected route, * - FIB route,p - stale info
Timers: Uptime

IP Route Table for VRF "default"
S   *> ::/0 [1/0] via 6000::2, xe5, 00:00:18
C   *> ::1/128 via ::, lo, 20:46:46
C   *> 3333::3/128 via ::, lo, 00:01:08
C   *> 6000::/64 via ::, xe5, 00:00:57
C   *> fe80::/64 via ::, ce43, 01:50:18
C      fe80::/64 via ::, xe32, 01:50:18
C      fe80::/64 via ::, xe29, 01:50:18
C      fe80::/64 via ::, xe15, 01:50:18
C      fe80::/64 via ::, xe3, 01:50:18
C      fe80::/64 via ::, ce47, 04:01:18
C      fe80::/64 via ::, ce45, 04:01:18
C      fe80::/64 via ::, xe42, 04:01:18
C      fe80::/64 via ::, xe41, 04:01:18
C      fe80::/64 via ::, xe5, 04:01:18
C      fe80::/64 via ::, xe34, 20:46:15
C      fe80::/64 via ::, xe33, 20:46:15
C      fe80::/64 via ::, xe12, 20:46:15
C      fe80::/64 via ::, xe11, 20:46:15
```

© 2026 IP Infusion Inc. Proprietary

Static Route Discard Configuration

This section 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 7 shows the configuration required to enable static route discard for IPv4.

Topology

Figure 7. Static route discard topology

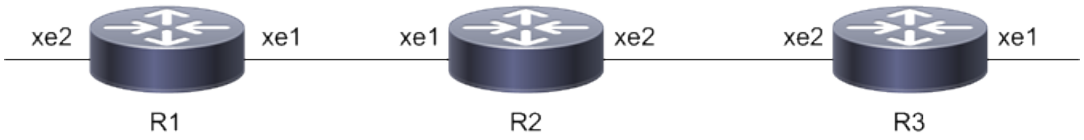

Configuration

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

© 2026 IP Infusion Inc. Proprietary

`R1(config-router)#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)#ip prefix-list p1`	Configure IP prefix list
`R2(config-ip-prefix-list)#seq 5 permit any`	Create an access rule to permit any IP packets
`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)#redistribute connected`	Advertise the connected network into BGP
`R2(config-router)#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)#end`	Exit BGP router mode

© 2026 IP Infusion Inc. Proprietary

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
`R3(config-router)#redistribute connected`	Advertise the connected network into BGP
`R3(config-if)#commit`	Commit the candidate configuration to the running configuration
`R3(config-router)#end`	Exit BGP router.

Validation

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

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
 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 7 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

© 2026 IP Infusion Inc. Proprietary

`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-router-af)#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`	Configure IPv6 prefix list.
`R2(config-ipv6-prefix-list)#permit any`	Create an access rule to permit any IP packets
`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-af)#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

© 2026 IP Infusion Inc. Proprietary

`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-af)#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
```

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

The figure below represents a sample network topology featuring RTR1, RTR2, RTR3, and RTR4, used to demonstrate the static route object tracking feature.

Figure 8. Static Route Object Tracking

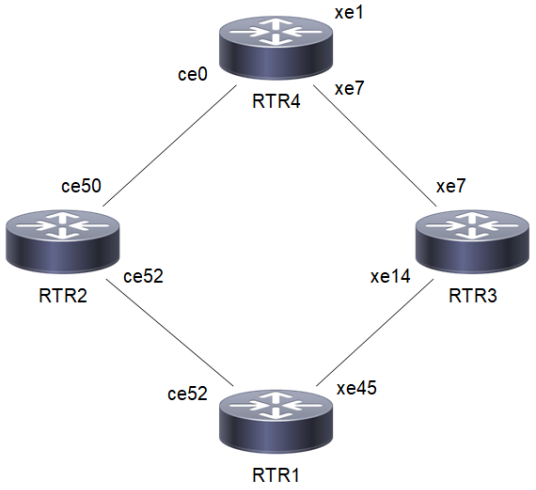

Static Route Object Tracking IPv4 Configuration

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.

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

`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.

Static Route Object Tracking IPv4 Validation

RTR1: IP SLA Routing and Tracking Details

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: Remove IPv4 Address

```
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: Down State in IP SLA Routing

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: Add IPv4 Address

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: UP State in IP SLA Routing and Tracking

By configuring the IPv4 address on `RTR2interfacece52` 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
```

Static Route Object Tracking 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.

© 2026 IP Infusion Inc. Proprietary

`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.

© 2026 IP Infusion Inc. Proprietary

`RTR4(config)#exit`	Exit configure mode.

Static Route Object Tracking IPv6 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
```

© 2026 IP Infusion Inc. Proprietary

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.

Configuration

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

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.

© 2026 IP Infusion Inc. Proprietary

Topology

A series of configurations were implemented on routers R1, R2, and R3, as well as on the Router A, 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 9. Route Monitor Topology

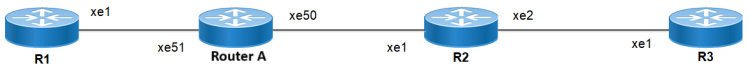

IPv4 Route Monitor Configuration

Router A Configuration

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.

`Router A#configure terminal`	Enter configure mode.
`Router A(config)#interface xe50`	Enter interface mode `xe50`.
`Router A(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`.
`Router A(config-if)#exit`	Exit interface mode `xe50`.
`Router A(config)#interface xe51`	Enter interface mode `xe51`.
`Router A(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.
`Router A(config-if)#exit`	Exit interface mode `xe51`.
`Router A(config)#ip sla 1`	Create an IP SLA operation with index `1`.
`Router A(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.
`Router A(config-ip-sla-echo)#threshold 1000`	Set the threshold value for SLA to `1000` milliseconds.
`Router A(config-ip-sla-echo)#timeout 1000`	Set the timeout value for SLA to `1000` milliseconds.
`Router A(config-ip-sla-echo)#frequency 5`	Configure the frequency value for SLA to send ICMP echo packets every 5 seconds.
`Router A(config-ip-sla-echo)#exit`	Exit IP SLA echo mode.
`Router A(config-ip-sla)#exit`	Exit IP SLA mode.
`Router A(config)#time-range tr1`	Create a time range named `tr1`.
`Router A(config-tr)#start-time 11:22 3 july 2021`	Set the start time for the time range to `11:22 on`

	`July 3, 2021`.
`Router A(config-tr)#end-time after 200`	Set the end time to be `200` minutes from the start time.
`Router A(config-tr)#exit`	Exit time-range mode.
`Router A(config)#ip sla schedule 1 time-range tr1`	Schedule IP SLA operation 1 to run within the specified time range `tr1`.
`Router A(config)#track 1 ip sla 1 reachability`	Creating a tracking object to monitor the reachability status of IP SLA operation `1`.
`Router A(config-object-track)#exit`	Exit object track mode.
`Router A(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`.
`Router A(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`.
`Router A(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`.
`Router A(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`.
`Router A(config)#commit`	Commit the candidate configuration to the running configuration.
`Router A(config)#interface xe51`	Enter interface mode `xe51`.
`Router A(config-if)#object-tracking all`	Enable object tracking for `all` tracking objects on interface `xe51`.
`Router A(config-if)#object-tracking 1`	Configure object tracking `1` on interface `xe51`.
`Router A(config-if)#object-tracking 2`	Configure object tracking `2` on interface `xe51`.
`Router A(config-if)#exit`	Exit interface mode.
`Router A(config)#exit`	Exit configure mode.

Routing and Static Routes Configuration

The following commands configure routers R1, R2, and R3 to forward network traffic to their designated destinations within the network. They assign IP addresses to the interfaces connected to the routing devices and add static routes for destination networks with a next-hop IP.

R1

```
!
interface xe1
 ip address 1.1.1.1/24
commit
!
ip route 2.2.2.0/24 1.1.1.2
ip route 3.3.3.0/24 1.1.1.2
commit
!
```

R2

```
!
interface xe1
 ip address 2.2.2.2/24
commit
!
interface xe2
 ip address 3.3.3.1/24
commit
!
ip route 1.1.1.0/24 2.2.2.1
commit
!
```

R3

```
!
interface xe1
 ip address 3.3.3.2/24
commit
!
ip route 1.1.1.0/24 3.3.3.1
ip route 2.2.2.0/24 3.3.3.1
commit
!
```

IPv4 Route Monitor 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.

Router A

```
Router A#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

Router A#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]

Router A#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
```

© 2026 IP Infusion Inc. Proprietary

```
Router A#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 Route Monitor Configuration

Router A Configuration

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.

`Router A#configure terminal`	Enter configure mode.
`Router A(config)#interface xe50`	Enter interface mode `xe50`.
`Router A(config-if)#ipv6 address 2000::1/64`	Assign an IPv6 address (`2000::1/64`) to interface `xe50`.
`Router A(config-if)#exit`	Exit interface mode `xe50`.
`Router A(config)#interface xe51`	Enter interface mode `xe51`.
`Router A(config-if)#ipv6 address 1000::2/64`	Assign an IPv6 address (`1000::2/64`) to interface `xe51`.
`Router A(config-if)#exit`	Exit interface mode `xe51`.
`Router A(config)#ip sla 1`	Create an IP SLA operation with index `1`.
`Router A(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.
`Router A(config-ip-sla-echo)#threshold 1000`	Set the threshold value for SLA to `1000` milliseconds.
`Router A(config-ip-sla-echo)#timeout 1000`	Set the timeout value for SLA to `1000` milliseconds.
`Router A(config-ip-sla-echo)#frequency 5`	Configure the frequency value for SLA to send IPv6 ICMP echo packets every 5 seconds.
`Router A(config-ip-sla-echo)#exit`	Exit IP SLA echo mode.
`Router A(config-ip-sla)#exit`	Exit IP SLA mode.

© 2026 IP Infusion Inc. Proprietary

`Router A(config)#time-range tr1`	Create a time range named `tr1`.
`Router A(config-tr)#start-time 11:22 3 july 2021`	Set the start time for the time range to `11:22 on July 3, 2021`.
`Router A(config-tr)#end-time after 200`	Set the end time to be `200` minutes from the start time.
`Router A(config-tr)#exit`	Exit time-range mode.
`Router A(config)#ip sla schedule 1 time-range tr1`	Schedule IP SLA operation `1` to run within the specified time range `tr1`.
`Router A(config)#track 1 ip sla 1 reachability`	Creating a tracking object to monitor the reachability status of IP SLA operation `1`.
`Router A(config-object-track)#exit`	Exit object track mode.
`Router A(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`.
`Router A(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.
`Router A(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`.
`Router A(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`.
`Router A(config)#commit`	Commit the candidate configuration to the running configuration.
`Router A(config)#interface xe51`	Enter interface mode xe51.
`Router A(config-if)#object-tracking all`	Enable object tracking for `all` tracking objects on interface `xe51`.
`Router A(config-if)#object-tracking 1`	Configure object tracking `1` on interface `xe51`.
`Router A(config-if)#object-tracking 2`	Configure object tracking `2` on interface `xe51`.
`Router A(config-if)#exit`	Exit interface mode.
`Router A(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 Configuration

`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 Configuration

`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 Configuration

`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.

IPv6 Route Monitor 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.

Router A Validation

```
Router A#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

Router A#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]

Router A#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
------------------------------------------------------------------------

Router A#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

Router A#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.

Route Monitor Command

The Route Monitor feature introduces the object-tracking (page 332) configuration command. For more information, refer to the IP Service Level Agreements Commands (page 318), and Object Tracking Commands (page 330) and *Interface Commands* sections in the *OcNOS System Management* document.

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.

© 2026 IP Infusion Inc. Proprietary

Default

None

Command Mode

Interface mode

Applicability

Introduced in OcNOS version 6.4.1.

Example

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

```
OcNOS(config)#interface 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.

Glossary

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

Key Terms/Acronym	Description
Acronym	Description
NSM	Network and Service Management
IP SLA	Internet Protocol Service Level Agreement
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

© 2026 IP Infusion Inc. Proprietary

Key Terms/Acronym	Description
	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.

Hybrid Switch Router Configuration

This section describes configurations that can be done with the Hybrid Switch Router solution. OcNOS can be configured as a Layer 2 switch, a Layer 3 router or a Hybrid Switch Router.

Overview

An OcNOS Hybrid Switch Router offers Layer 3 forwarding found in routers with the high-speed performance associated with traditional Layer 2 switches. The following are some advantages of Hybrid Switch Routers:

- Reduced system cost and infrastructure. Traditionally you would require a separate box for switching and one for routing.
- Off-loading IP traffic from backbone routers, thus making them more efficient for firewalls and WAN connectivity.
- Simplified network design and maintenance.

Routing and Switching

Layer 2 and Layer 3 switches are similar at a high-level, both look at the packet headers, and steer the packets toward their destination port. Therefore, after being passed through a switch or router, the packet is closer to its destination.

Layer 2 Switching

Layer 2 switches are typically used to provide connectivity within high bandwidth local area networks (LANs). A Layer 2 switch makes forwarding decisions based on the MAC or the Layer 2 header. It extracts the Layer 2 header from the packet, finds a matching destination address in the forwarding table, and transmits the packet out to the port associated with the specific destination address in the forwarding table. The forwarding table is populated through a self-learning process, whereby each arriving packet is used to update the entries in the table. Typically, the Layer 2 switch implements the switching function in the hardware, as that requires stripping of the packet only in two layers (the physical and data link layer) to get to the useful part of the packet header. This allows switches to steer packets at wire-speed rates without slowing down arriving streams of packets to process them.

Layer 3 Routing

Layer 3 (L3) routers are typically used to provide connectivity between different LANs. A Layer 3 router discards MAC headers, and indexes further into the packet–making decisions based on the IP or Layer 3 header. It extracts the Layer-3 header from the packet, finds a matching destination address in a routing table, identifies a new MAC address for the packet from an ARP cache, wraps the IP packet in a new MAC header, and then transmits the packet out to the port associated with that destination address in the routing table.

The routing table is populated through statically configured command line interface entries or through routing protocol messages from neighboring routers. A Layer 3 router must strip through 3 layers (physical, data link, and network) which is more complicated than a Layer 2 switch. Layer 3 routers historically implement the routing function in software. This often results in limited packet-forwarding rates. However, improvements in VLSI circuit technology have allowed Layer 3 routing functions to be implemented rapidly in hardware, enabling wire-speed performance similar to the performance of Layer 2 switches. As a result, along with the complexity of next-generation Layer 3 routers, the throughput of these routers has also been increasing.

© 2026 IP Infusion Inc. Proprietary

An architecture is required that is flexible enough to accommodate the demands of different customers, and accommodate the changing demands of a single customer whose requirements may change over time. Typical Layer 2 switches and Layer 3 routers fail to provide this flexibility.

An optimal configuration can be an integrated solution, a Layer 3 router with Layer 2 bridge groups around it. The OcNOS Hybrid Switch Router implementation allows easy configuration of different combinations of routers and switches. OcNOS can be configured as an absolute Layer 3 router, absolute Layer 2 switch (Figure 10) or a hybrid Layer 2/Layer 3 switch router, (Figure 11) that can easily change modes with the use of a single command.

System Configuration

- OcNOS stack will bring up all the ports of the board as routed ports by default.
- However OcNOS provides flexibility to create a Layer 2 bridge, and ports can be converted to switch ports and added to the bridge.
- OcNOS also supports VLAN interfaces and routing between VLANs.

Thus it can work as a router, a switch, or as a hybrid switch.

Hybrid Switch Router Possibilities

With only Layer 2 protocols configured, the OcNOS Hybrid Switch Router can become an absolute Layer 2 switch.

Figure 10. Layer 2 Switch

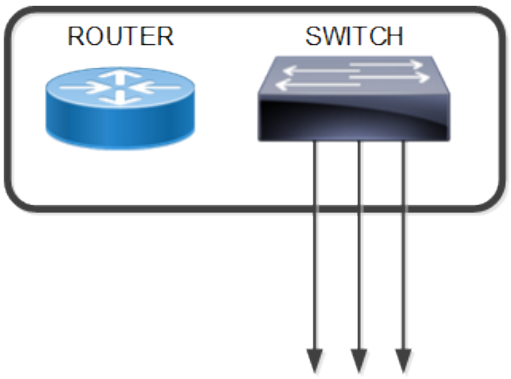

With Layer 2 and Layer 3 protocols configured, the OcNOS Hybrid Switch Router can become a Switch and/or a Router.

Figure 11. Working as a Router or a Switch

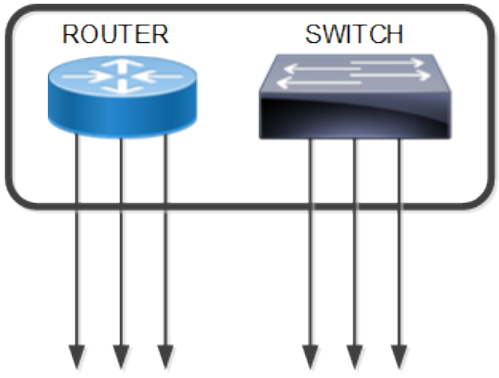

With only Layer 3 protocols enabled, the OcNOS Hybrid Switch Router can become an absolute router.

© 2026 IP Infusion Inc. Proprietary

Figure 12. Working as Layer 3 Router

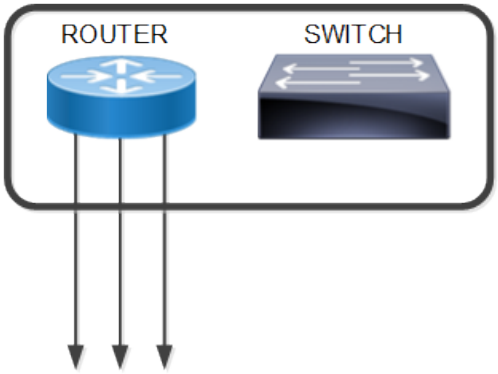

On switch ports, VLANs can be created for different broadcast domains.

Figure 13. Port- or Policy-based VLANs

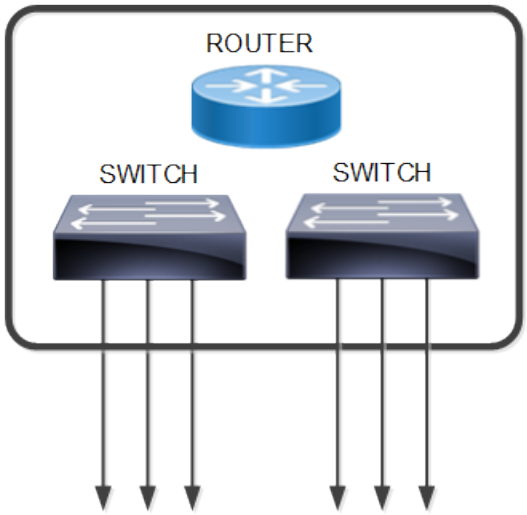

For routing between VLANs, the OcNOS routing protocols or static routing via NSM can be utilized.

Figure 14. Routing between VLANs

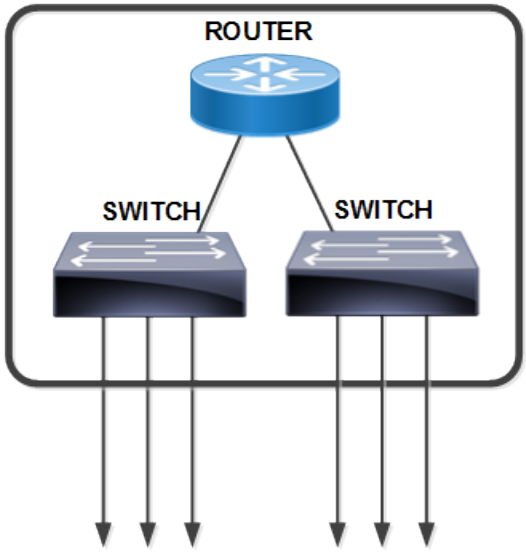

© 2026 IP Infusion Inc. Proprietary

For routing between VLANs and other routing ports, OcNOS routing protocols or static routing via NSM can be utilized.

Figure 15. Routing between VLANs and Routing Ports

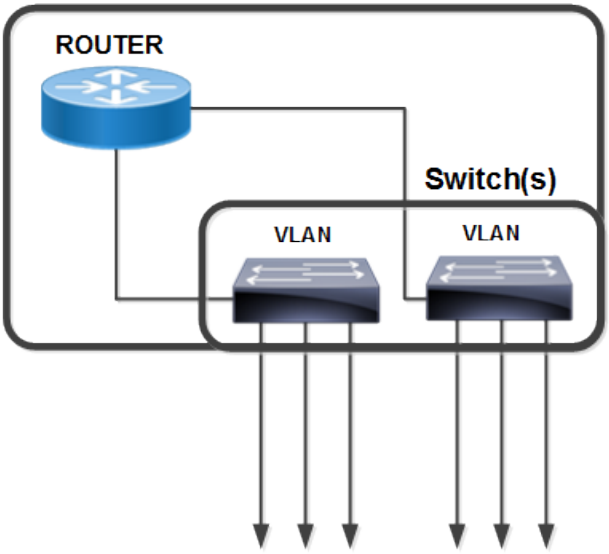

Configuring Layer 2 Interfaces

For the Hybrid Switch Router, it is important to understand that by default, all interfaces are configured as routed interfaces. To configure a Layer 2 interface (switched interface), you must explicitly configure this using the `switchport` command in the interface mode. For example:

`#configure terminal`	Enter the Configure mode.
`(config)#bridge 1 protocol mstp`	Create a MSTP bridge.
`(config)#interface eth2`	Specify an interface to configure and enter the Interface mode.
`(config-if)#switchport`	Configure eth2 as a Layer 2 port.
`(config-if)#bridge-group 1`	Associate interface to bridge 1.
`(config-if)#no shutdown`	Start interface.eth2
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode

Configuring Layer 3 Interfaces

For the Hybrid Switch Router, it is important to understand that by default, all interfaces in OcNOS are L3 ports. If a port has been configured as L2, then use the `no switchport` command to make it a L3 port.

For example:

`#configure terminal`	Enter the Configure mode.

`(config)#interface eth3`	Specify an interface to configure and enter the Interface mode.
`(config-if)#no switchport`	Configure eth3 as a Layer 3 port.
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode

In the Hybrid Switch Router mode, if a VLAN is configured, a Layer 3 interface based on the bridge-group number and VLAN ID is created. This Layer 3 interface is advertised to all the Layer 3 protocols. For example:

`#configure terminal`	Enter the Configure mode.
`(config)#bridge 1 protocol mstp`	Creating bridge.
`(config)#vlan database`	Enter the VLAN configure mode.
`(config-vlan)#vlan 2 bridge 1`	Enable VLAN 2 on bridge 1.
`(config-vlan)#commit`	Commit the candidate configuration to the running configuration
`(config-vlan)#exit`	Exit VLAN configure mode

The last step in the example above show an interface given a name with the following format:

```
vlanXX.YYYY
```

Where `XX` is the bridge ID, and `YYYY` is the VLAN ID,

For example, the name, `vlan1.3` indicates that VLAN IP's interface is in VLAN 3, and bridge-group 1.

© 2026 IP Infusion Inc. Proprietary

Policy Based Routing Configuration

This section contains a sample Policy Based Routing (PBR) configuration.

Policy Based Routing (PBR) is an advanced packet forwarding feature which is different from conventional destination address based routing. Policy Based Routing (PBR) allows data packets forwarding based on policies defined by network administrators.

In conventional routing, when a packet is received on the router, destination address in the packet is looked upon in the routing table and if the routing entry is found, packet is routed based on routing entry. In policy based routing, routing decision could be made from source address, destination address, transport protocol id, source port, destination port, or a combination of these criteria.

PBR includes a mechanism for selectively applying policies based on an access list or other criteria. Actions taken might include (a) Forwarding a packet to a directly connected ip nexthop (b) Black hole/Drop. If traffic doesn't match the route-map's match criteria, then it will be routed as if no PBR policy exists. PBR config is interface oriented, hence when applied it affects only the traffic ingressing on that interface. It does not apply on traffic egressing on that interface or traffic ingressing on an interface without a pbr route-map.

Topology

Figure 16. PBR Routing Toplogy

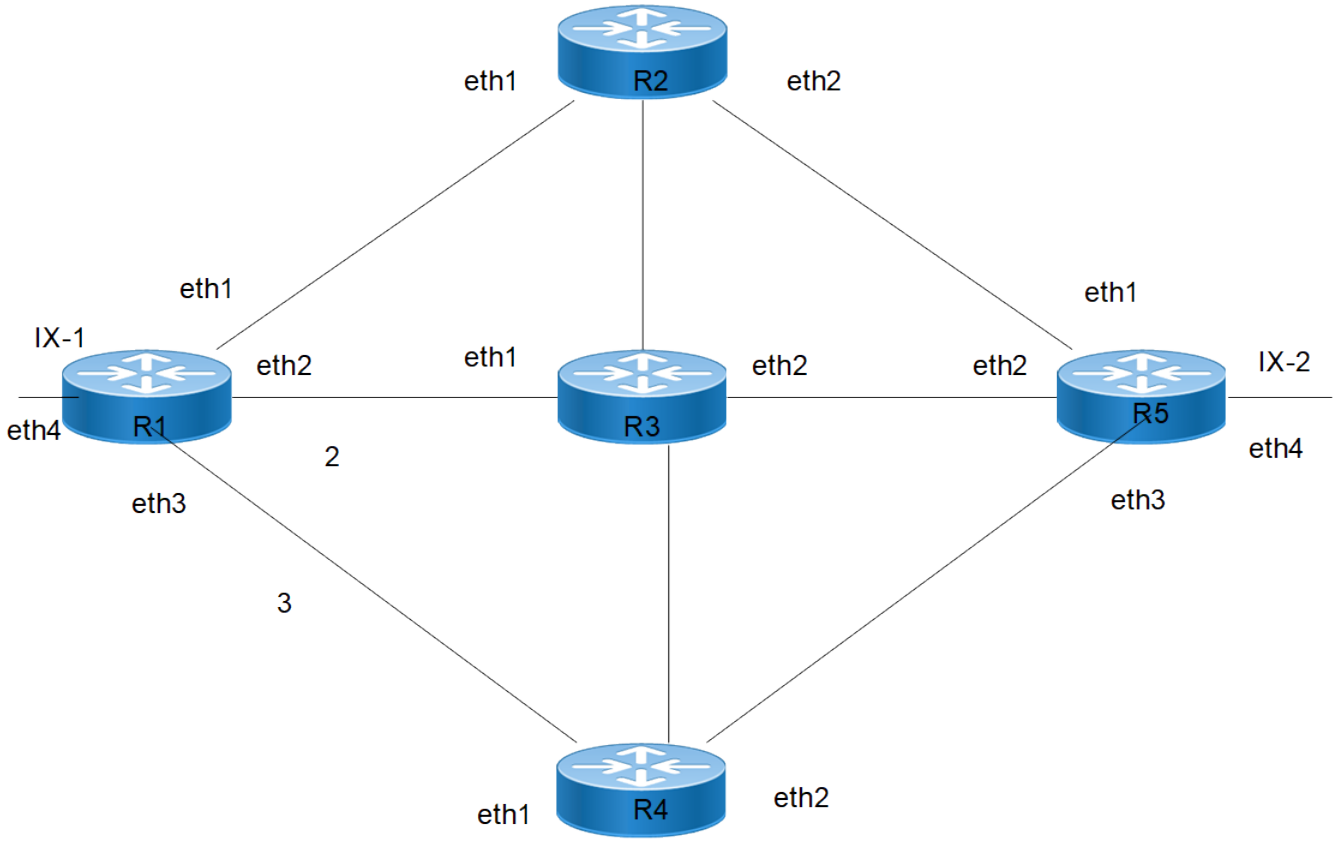

Configurations - PBR IPv4

R1

`R1#configure terminal`	Enter configure mode.
`R1(config)#feature pbr`	Enable PBR support
`R1(config)#hardware-profile filter ingress- ipv4-ext enable`	Enable Ingress IPv4 group extended for PBR support
`R1(config)#ip access-list 123`	Create ip access-list named 123
`R1(config-ip-acl)#10 permit any 101.1.1.0/24 201.1.1.0/24`	Create an access rule to permit IP packets with source 101.1.1.0/24 and destination 201.1.1.0/24
`R1(config-ip-acl)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-ip-acl)#exit`	Exit access-list mode
`R1(config)#route-map 123 permit 10`	Configure route-map with name 123 and sequence number 10
`R1(config-route-map)#match ip address 123`	Match ip address with ACL 123
`R1(config-route-map)#set ip next-hop 13.1.1.2`	Set next-hop to forward the matching IP packets
`R1(config-route-map)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-route-map)#exit`	Exit route-map mode
`R1(config)#interface lo`	Enter interface mode.
`R1(config-if)#ip address 1.1.1.1/32`	Configure the IP address of the interface.
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#interface eth1`	Enter interface mode.
`R1(config-if)#ip address 2.1.1.1/24`	Configure the IP address of the interface.
`R1(config)#interface eth2`	Enter interface mode.
`R1(config-if)#ip address 6.1.1.1/24`	Configure the IP address of the interface.
`R1(config-if)#ip ospf cost 2`	Configuring ospf cost as 2
`R1(config)#interface eth3`	Enter interface mode.
`R1(config-if)#ip address 13.1.1.1/24`	Configure the IP address of the interface.
`R1(config-if)#ip ospf cost 3`	Configuring ospf cost as 3
`R1(config)#interface eth4`	Enter interface mode.
`R1(config-if)#ip address 101.1.1.2/24`	Configure the IP address of the interface.
`R1(config-if)#ip policy route-map 123`	Attach PBR on the ingress interface
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#router ospf 1`	Set the routing process ID .
`R1(config-router)# ospf router-id 1.1.1.1`	Configure OSPF router-id

© 2026 IP Infusion Inc. Proprietary

`R1(config-router)#network 1.1.1.1/32 area 0.0.0.0`	Configure OSPF network in area 0
`R1(config-router)# network 2.1.1.0/24 area 0.0.0.0`	Configure OSPF network in area 0
`R1(config-router)#network 6.1.1.0/24 area 0.0.0.0`	Configure OSPF network in area 0
`R1(config-router)#network 13.1.1.0/24 area 0.0.0.0`	Configure OSPF network in area 0
`R1(config-router)#network 101.1.1.0/24 area 0.0.0.0`	Configure OSPF network in area 0
`R1(config-router)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-router)#exit`	Exit router mode.

R2

`R2#configure terminal`	Enter configure mode.
`R2(config)#interface lo`	Enter interface mode.
`R2(config-if)#ip address 2.2.2.2/32`	Configure the IP address of the interface.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#interface eth1`	Enter interface mode.
`R2(config-if)#ip address 2.1.1.2/24`	Configure the IP address of the interface.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#interface eth2`	Enter interface mode.
`R2(config-if)#ip address 10.1.1.1/24`	Configure the IP address of the interface.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#router ospf 1`	Set the routing process ID .
`R2(config-router)# ospf router-id 2.2.2.2`	Configure OSPF router-id
`R2(config-router)#network 2.2.2.2/32 area 0.0.0.0`	Configure OSPF network in area 0
`R2(config-router)# network 2.1.1.0/24 area 0.0.0.0`	Configure OSPF network in area 0
`R2(config-router)#network 10.1.1.0/24 area 0.0.0.0`	Configure OSPF network in area 0
`R2(config-router)#commit`	Commit the candidate configuration to the running configuration.
`R2(config-router)#exit`	Exit router mode.

R3

`R3#configure terminal`	Enter configure mode.
`R3(config)#interface lo`	Enter interface mode.
`R3(config-if)#ip address 3.3.3.3/32`	Configure the IP address of the interface.
`R3(config-if)#exit`	Exit interface mode.
`R3(config)#interface eth1`	Enter interface mode.
`R3(config-if)#ip address 6.1.1.2/24`	Configure the IP address of the interface.

R3(config-if)#exit	Exit interface mode.
R3(config)#interface eth2	Enter interface mode.
R3(config-if)#ip address 12.1.1.1/24	Configure the IP address of the interface.
R3(config-if)#exit	Exit interface mode.
R3(config)#router ospf 1	Set the routing process ID .
R3(config-router)# ospf router-id 3.3.3.3	Configure OSPF router-id
R3(config-router)#network 3.3.3.3/32 area 0.0.0.0	Configure OSPF network in area 0
R3(config-router)# network 6.1.1.0/24 area 0.0.0.0	Configure OSPF network in area 0
R3(config-router)#network 12.1.1.0/24 area 0.0.0.0	Configure OSPF network in area 0
R3(config-router)#commit	Commit the candidate configuration to the running configuration.
R3(config-router)#exit	Exit router mode.

R4

R4#configure terminal	Enter configure mode.
R4(config)#interface lo	Enter interface mode.
R4(config-if)#ip address 4.4.4.4/32	Configure the IP address of the interface.
R4(config-if)#exit	Exit interface mode.
R4(config)#interface eth1	Enter interface mode.
R4(config-if)#ip address 13.1.1.2/24	Configure the IP address of the interface.
R4(config-if)#exit	Exit interface mode.
R4(config)#interface eth2	Enter interface mode.
R4(config-if)#ip address 15.1.1.1/24	Configure the IP address of the interface.
R4(config-if)#exit	Exit interface mode.
R4(config)#router ospf 1	Set the routing process ID .
R4(config-router)# ospf router-id 4.4.4.4	Configure OSPF router-id
R4(config-router)#network 4.4.4.4/32 area 0.0.0.0	Configure OSPF network in area 0
R4(config-router)# network 13.1.1.0/24 area 0.0.0.0	Configure OSPF network in area 0
R4(config-router)#network 15.1.1.0/24 area 0.0.0.0	Configure OSPF network in area 0
R4(config-router)#commit	Commit the candidate configuration to the running configuration.
R4(config-router)#exit	Exit router mode.

R5

R5#configure terminal	Enter configure mode.
R5(config)#interface lo	Enter interface mode.

R5(config-if)#ip address 5.5.5.5/32	Configure the IP address of the interface.
R5(config-if)#exit	Exit interface mode.
R5(config)#interface eth1	Enter interface mode.
R5(config-if)#ip address 10.1.1.2/24	Configure the IP address of the interface.
R4(config-if)#exit	Exit interface mode.
R5(config)#interface eth2	Enter interface mode.
R5(config-if)#ip address 12.1.1.1/24	Configure the IP address of the interface.
R5(config-if)#exit	Exit interface mode.
R5(config)#interface eth3	Enter interface mode.
R5(config-if)#ip address 15.1.1.2/24	Configure the IP address of the interface.
R5(config-if)#exit	Exit interface mode.
R5(config)#interface eth4	Enter interface mode.
R5(config-if)#ip address 202.1.1.2/24	Configure the IP address of the interface.
R5(config-if)#exit	Exit interface mode.
R5(config)#router ospf 1	Set the routing process ID .
R5(config-router)# ospf router-id 5.5.5.5	Configure OSPF router-id
R5(config-router)#network 5.5.5.5/32 area 0.0.0.0	Configure OSPF network in area 0
R5(config-router)# network 10.1.1.0/24 area 0.0.0.0	Configure OSPF network in area 0
R5(config-router)# network 12.1.1.0/24 area 0.0.0.0	Configure OSPF network in area 0
R5(config-router)#network 15.1.1.0/24 area 0.0.0.0	Configure OSPF network in area 0
R5(config-router)#network 202.1.1.0/24 area 0.0.0.0	Configure OSPF network in area 0
R5(config-router)#commit	Commit the candidate configuration to the running configuration.
R5(config-router)#exit	Exit router mode.

Validation

R1

```
R1#show  run aclmgr
ip access-list 123
 10 permit any 101.1.1.0/24 201.1.1.0/24
R1#show  run interface eth4
!
interface eth4
 ip address 101.1.1.2/24
 ip policy route-map 123
!
R1#
R1#show  route-map 123
route-map 123, permit, sequence 10
  Match clauses:
    ip address: 123
  Set clauses:
    ip next-hop 13.1.1.2
R1#
R1#show  ip policy
```

```
Interface      Route-map               Status    VRF-Name
  eth4            123                        Active           default

R1#
R1#clear route-map 123 pbr-statistics
R1#show route-map 123 pbr-statistics

Route-map 123, family IP
IP PBR Count: 1
VRF-name: default
  Sequence 10, permit
    Policy routing matches: 38764427 packets
    Current action in HW: Route
```

Configuration - PBR IPv6

R1

`R1#configure terminal`	Enter configure mode.
`R1(config)#feature pbr`	Enable PBR support
`R1(config)#hardware-profile filter ingress- ipv6 enable`	Enable Ingress IPv6 group for PBR support
`R1(config)# ipv6 access-list 123`	Create ipv6 access-list named 123
`R1(config-ipv6-acl)#10 permit any 101::/64 202::/64`	Create an access rule to permit IPv6 packets with source 101::/64 and destination 202::/64
`R1(config-ipv6-acl)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-ipv6-acl)#exit`	Exit access-list mode
`R1(config)#route-map 123 permit 10`	Configure route-map with name 123 and sequence number 10
`R1(config-route-map)# match ipv6 address 123`	Match ip address with ACL 123
`R1(config-route-map)#set ipv6 next-hop 6111::2`	Set next-hop to forward the matching IP packets
`R1(config-route-map)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-route-map)#exit`	Exit route-map mode
`R1(config)#interface lo`	Enter interface mode.
`R1(config-if)#ip address 1.1.1.1/32`	Configure the IPv6 address of the interface.
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#router ipv6 ospf 100`	Creating OSPFv3 routing instance
`R1(config-router)#exit`	Exit router mode.
`R1(config)#interface eth1`	Enter interface mode.
`R1(config-if)#ipv6 address 2111::1/64`	Configure the IPv6 address of the interface.
`R1(config-if)#ipv6 router ospf area 0.0.0.0 tag 100`	Enable OSPFv3 routing on an interface, and assign the Area ID 0.

© 2026 IP Infusion Inc. Proprietary

R1(config-if)#exit	Exit interface mode.
R1(config)#interface eth2	Enter interface mode.
R1(config-if)#ipv6 address 6111::1/64	Configure the IPv6 address of the interface.
R1(config-if)#ipv6 router ospf area 0.0.0.0 tag 100	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
R1(config-if)# ipv6 ospf cost 2	Configuring ospf cost as 2
R1(config-if)#exit	Exit interface mode.
R1(config)#interface eth3	Enter interface mode.
R1(config-if)#ipv6 address 1311::1/64	Configure the IPv6 address of the interface.
R1(config-if)#ipv6 router ospf area 0.0.0.0 tag 100	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
R1(config-if)# ipv6 ospf cost 3	Configuring ospf cost as 3
R1(config-if)#exit	Exit interface mode.
R1(config)#interface eth4	Enter interface mode.
R1(config-if)#ipv6 router ospf area 0.0.0.0 tag 100	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
R1(config-if)#ipv6 address 101::2/64	Configure the IPv6 address of the interface.
R1(config-if)#ipv6 policy route-map 123	Attach PBR on the ingress interface
R1(config-if)#commit	Commit the candidate configuration to the running configuration.
R1(config-if)#exit	Exit interface mode.

R2

R2#configure terminal	Enter configure mode.
R2(config)#interface lo	Enter interface mode.
R2(config-if)#ip address 2.2.2.2/32	Configure the IPv6 address of the interface.
R2(config-if)#exit	Exit interface mode.
R2(config)#router ipv6 ospf 100	Creating OSPFv3 routing instance
R2(config-router)#exit	Exit router mode.
R2(config)#interface eth1	Enter interface mode.
R2(config-if)#ipv6 address 2111::2/64	Configure the IPv6 address of the interface.
R2(config-if)#ipv6 router ospf area 0.0.0.0 tag 100	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
R2(config-if)#exit	Exit interface mode.
R2(config)#interface eth2	Enter interface mode.
R2(config-if)#ipv6 address 1011::1/64	Configure the IPv6 address of the interface.
R2(config-if)#ipv6 router ospf area 0.0.0.0 tag 100	Enable OSPFv3 routing on an interface, and assign the Area ID 0.

© 2026 IP Infusion Inc. Proprietary

R2(config-if)#commit	Commit the candidate configuration to the running configuration.
R2(config-if)#exit	Exit interface mode.

R3

R3#configure terminal	Enter configure mode.
R3(config)#router ipv6 ospf 100	Creating OSPFv3 routing instance
R3(config-router)#exit	Exit router mode.
R3(config)#interface lo	Enter interface mode.
R3(config-if)#ip address 3.3.3.3/32	Configure the IPv6 address of the interface.
R3(config-if)#exit	Exit interface mode.
R3(config)#interface eth1	Enter interface mode.
R3(config-if)#ipv6 address 6111::2/64	Configure the IPv6 address of the interface.
R3(config-if)#ipv6 router ospf area 0.0.0.0 tag 100	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
R2(config-router)#exit	Exit router mode.
R3(config)#interface eth2	Enter interface mode.
R3(config-if)#ipv6 address 1211::1/64	Configure the IPv6 address of the interface.
R3(config-if)#ipv6 router ospf area 0.0.0.0 tag 100	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
R2(config-if)#commit	Commit the candidate configuration to the running configuration.
R3(config-if)#exit	Exit interface mode.

R4

R4#configure terminal	Enter configure mode.
R4(config)#router ipv6 ospf 100	Creating OSPFv3 routing instance
R4(config-router)#exit	Exit router mode.
R4(config)#interface lo	Enter interface mode.
R4(config-if)#ip address 4.4.4.4/32	Configure the IPv6 address of the interface.
R4(config-if)#exit	Exit interface mode.
R4(config)#interface eth1	Enter interface mode.
R4(config-if)#ipv6 address 1311::2/64	Configure the IPv6 address of the interface.
R4(config-if)#ipv6 router ospf area 0.0.0.0 tag 100	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
R4(config-if)#exit	Exit interface mode.
R4(config)#interface eth2	Enter interface mode.

© 2026 IP Infusion Inc. Proprietary

`R4(config-if)#ipv6 address 1511::1/64`	Configure the IPv6 address of the interface.
`R4(config-if)#ipv6 router ospf area 0.0.0.0 tag 100`	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
`R4(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R4(config-if)#exit`	Exit interface mode.

R5

`R5#configure terminal`	Enter configure mode.
`R5(config)#router ipv6 ospf 100`	Creating OSPFv3 routing instance
`R5(config-router)#exit`	Exit router mode.
`R5(config)#interface lo`	Enter interface mode.
`R5(config-if)#ip address 5.5.5.5/32`	Configure the IPv6 address of the interface.
`R5(config-if)#exit`	Exit interface mode.
`R5(config)#interface eth1`	Enter interface mode.
`R5(config-if)#ipv6 address 1011::2/64`	Configure the IPv6 address of the interface.
`R5(config-if)#ipv6 router ospf area 0.0.0.0 tag 100`	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
`R5(config-if)#exit`	Exit interface mode.
`R5(config)#interface eth2`	Enter interface mode.
`R5(config-if)#ipv6 address 1211::2/64`	Configure the IPv6 address of the interface.
`R5(config-if)#ipv6 router ospf area 0.0.0.0 tag 100`	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
`R5(config-if)#exit`	Exit interface mode.
`R5(config)#interface eth3`	Enter interface mode.
`R5(config-if)#ipv6 address 1511::2/64`	Configure the IPv6 address of the interface.
`R5(config-if)#ipv6 router ospf area 0.0.0.0 tag 100`	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
`R5(config-if)#exit`	Exit interface mode.
`R5(config)#interface eth4`	Enter interface mode.
`R5(config-if)#ipv6 address 202::2/64`	Configure the IPv6 address of the interface.
`R5(config-if)#ipv6 router ospf area 0.0.0.0 tag 100`	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
`R5(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R5(config-if)#exit`	Exit interface mode.

© 2026 IP Infusion Inc. Proprietary

Validation

R1

```
R1#show  run aclmgr
ipv6 access-list 123
 10 permit any 101::/64 202::/64
 268435453 permit icmpv6 any any
R1#show  run interface eth4
!
interface eth4
 ip address 101.1.1.2/24
 ipv6 address 101::2/64
 ipv6 policy route-map 123
 ipv6 router ospf area 0.0.0.0 tag 100
!
R1#show  route-map 123
route-map 123, permit, sequence 10
  Match clauses:
    ipv6 address: 123
  Set clauses:
    ipv6 next-hop 6111::2

R1#show  ipv6 policy
Interface      Route-map               Status    VRF-Name
    eth4            123                              Active    default

R1#
R1#clear route-map 123 pbr-statistics
R1#sho  route-map 123 pbr-statistics

Route-map 123, family IPv6
IPv6 PBR Count: 1
VRF-name: default
  Sequence 10, permit
    Policy routing matches: 1077577 packets
    Current action in HW: Route
```

PBR with Nexthop Tracking

Policy-Based Routing (PBR) with Next-Hop Tracking is a networking feature designed to increase routing resiliency and efficiency by dynamically tracking the availability of the next-hop IP address associated with a PBR policy. When a next-hop IP address becomes unreachable, the system typically waits for the ARP entry to expire before switching traffic to an alternative route, resulting in dropped traffic during this delay. This feature addresses the issue by enabling faster traffic redirection, minimizing disruptions and improving network reliability.

Policy-Based Routing (PBR) can be configured with a tracker-enabled next hop by using the optional tracking argument in the set nexthop command. When the tracker's state changes to UP, the tracker-enabled next hop is installed in hardware. If the tracker's state goes DOWN, the tracker-enabled next hop is removed—provided it is currently active—and PBR automatically reroutes traffic based on the policy’s fallback configuration.

It is important to note that these dynamic state changes do not affect the output of the show running-config command, which remains static regardless of the tracker's status.

Modify the R1 configuration to enable nexthop tracking. All other configurations will remain unchanged.

IPv4 Configurations for PBR with Nexthop Tracking

R1

`#configure terminal Enter`	configure mode.
`R1(config)#feature pbr`	Enable PBR support
`R1(config)#hardware-profile filter ingress- ipv4-ext enable`	Enable Ingress IPv4 group extended for PBR support
`R1(config)#ip access-list 123`	Create ip access-list named 123
`R1(config-ip-acl)#10 permit any 101.1.1.0/24 201.1.1.0/24`	Create an access rule to permit IP packets with source 101.1.1.0/24 and destination 201.1.1.0/24
`R1(config-ip-acl)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-ip-acl)#exit`	Exit access-list mode
`R1(config)#time-range tr1`	Configure a time-range.
`R1(config-tr)#start-time 11:22 3 july 2025`	Configure a start-time.
`R1(config-tr)#end-time after 200`	Configure end-time.
`R1(config-tr)#exit`	Exit time-range mode.
`R1(config)#ip sla 7`	Configure IP SLA with a unique number.
`R1(config-ip-sla)#icmp-echo ipv4 13.1.1.2 source-interface eth3`	Configure the ICMP-echo using destination IPv4 address and source interface name.
`R1(config-ip-sla-echo)#threshold 1000`	Configure the threshold value.
`R1(config-ip-sla-echo)#timeout 1000`	Configure the timeout value.
`R1(config-ip-sla-echo)#frequency 5`	Configure the frequency value.
`R1(config-ip-sla-echo)#exit`	Exit IP SLA echo mode.

© 2026 IP Infusion Inc. Proprietary

`R1(config-ip-sla)#exit`	Exit IP SLA mode.
`R1(config)#ip sla schedule 7 time-range tr1`	Schedule an IP SLA measurement.
`R1(config)#track 7 ip sla 7 reachability`	Configure the track ID for IP SLA.
`R1(config-object-track)#exit`	Exit object track mode.
`R1(config)#route-map 123 permit 10`	Configure route-map with name 123 and sequence number 10
`R1(config-route-map)#match ip address 123`	Match ip address with ACL 123
`R1(config-route-map)#set ip next-hop 13.1.1.2 verify-availability track 7`	Set next-hop to forward the matching IP packets, enable tracking for nexthop 13.1.1.2
`R1(config-route-map)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-route-map)#exit`	Exit route-map mode
`R1(config)#interface lo`	Enter interface mode.
`R1(config-if)#ip address 1.1.1.1/32`	Configure the IP address of the interface.
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#interface eth1`	Enter interface mode.
`R1(config-if)#ip address 2.1.1.1/24`	Configure the IP address of the interface.
`R1(config)#interface eth2`	Enter interface mode.
`R1(config-if)#ip address 6.1.1.1/24`	Configure the IP address of the interface.
`R1(config-if)#ip ospf cost 2`	Configuring ospf cost as 2
`R1(config)#interface eth3`	Enter interface mode.
`R1(config-if)#ip address 13.1.1.1/24`	Configure the IP address of the interface.
`R1(config-if)#ip ospf cost 3`	Configuring ospf cost as 3
`R1(config)#interface eth4`	Enter interface mode.
`R1(config-if)#ip address 101.1.1.2/24`	Configure the IP address of the interface.
`R1(config-if)#ip policy route-map 123`	Attach PBR on the ingress interface
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#router ospf 1`	Set the routing process ID .
`R1(config-router)# ospf router-id 1.1.1.1`	Configure OSPF router-id
`R1(config-router)#network 1.1.1.1/32 area 0.0.0.0`	Configure OSPF network in area 0
`R1(config-router)# network 2.1.1.0/24 area 0.0.0.0`	Configure OSPF network in area 0
`R1(config-router)#network 6.1.1.0/24 area 0.0.0.0`	Configure OSPF network in area 0
`R1(config-router)#network 13.1.1.0/24 area 0.0.0.0`	Configure OSPF network in area 0
`R1(config-router)#network 101.1.1.0/24 area 0.0.0.0`	Configure OSPF network in area 0
`R1(config-router)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-router)#exit`	Exit router mode.

© 2026 IP Infusion Inc. Proprietary

Validation

R1

```
OcNOS#show track
TRACK Id: 7
  IP SLA 7 reachability
  Reachability is UP
    5 changes, last change : 2025 Jun 19 08:02:07
Track interface :

OcNOS#show track summary
Object Tracking Summary
ID       Type        Type-Identifier       State
-----------------------------------------------------
7      ip-sla             7               UP

OcNOS#show nsm obj-track-table
Object Track Table
Track Id Track Status Last Updated
7        Up           2025 Jun 19 08:02:07



OcNOS#show hsl pbr

PBR name : 123, FIB ID : 0, Family : IP
Rule count : 1, Interface count : 1

Interface(RuleCnt) : eth4(2)

PBR rule number: 10
  MATCH details:
    IP-LIST   - ACL-name : 123 (IfRefCnt:2)
  SET details:
    Action : Route
    Next-hop count : 1
      Next-hop 13.1.1.2 (eth1), VALID, INSTALLED, NH 0xa40f49d8, TRACKED
    Load balance : NO
    Egress object ID : 0x20000c06
```

IPv6 Configurations for PBR with Nexthop Tracking

R1

`R1#configure terminal`	Enter configure mode.
`R1(config)#feature pbr`	Enable PBR support
`R1(config)#hardware-profile filter ingress- ipv6 enable`	Enable Ingress IPv6 group for PBR support
`R1(config)# ipv6 access-list 124`	Create ipv6 access-list named 1234
`R1(config-ipv6-acl)#10 permit any 101::/64 202::/64`	Create an access rule to permit IPv6 packets with source 101::/64 and destination 202::/64
`R1(config-ipv6-acl)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-ipv6-acl)#exit`	Exit access-list mode

© 2026 IP Infusion Inc. Proprietary

`R1(config)#time-range tr1`	Configure a time-range.
`R1(config-tr)#start-time 11:22 3 july 2025`	Configure a start-time.
`R1(config-tr)#end-time after 200`	Configure end-time.
`R1(config-tr)#exit`	Exit time-range mode.
`R1(config)#ip sla 8`	Configure IP SLA with a unique number.
`R1(config-ip-sla)#icmp-echo ipv6 1311::2 source-interface eth3`	Configure the ICMP-echo using destination IPv6 address and source interface name.
`R1(config-ip-sla-echo)#threshold 1000`	Configure the threshold value.
`R1(config-ip-sla-echo)#timeout 1000`	Configure the timeout value.
`R1(config-ip-sla-echo)#frequency 5`	Configure the frequency value.
`R1(config-ip-sla-echo)#exit`	Exit IP SLA echo mode.
`R1(config-ip-sla)#exit`	Exit IP SLA mode.
`R1(config)#ip sla schedule 8 time-range tr1`	Schedule an IP SLA measurement.
`R1(config)#track 8 ip sla 8 reachability`	Configure the track ID for IP SLA.
`R1(config-object-track)#exit`	Exit object track mode.
`R1(config)#route-map 124 permit 10`	Configure route-map with name 124 and sequence number 10
`R1(config-route-map)# match ipv6 address 124`	Match ip address with ACL 124
`R1(config-route-map)#set ipv6 next-hop 1311::2 verify-availability track 8`	Set next-hop to forward the matching IP packets
`R1(config-route-map)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-route-map)#exit`	Exit route-map mode
`R1(config)#interface lo`	Enter interface mode.
`R1(config-if)#ip address 1.1.1.1/32`	Configure the IPv4 address of the interface.
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#router ipv6 ospf 100`	Creating OSPFv3 routing instance
`R1(config-router)#exit`	Exit router mode.
`R1(config)#interface eth1`	Enter interface mode.
`R1(config-if)#ipv6 address 2111::1/64`	Configure the IPv6 address of the interface.
`R1(config-if)#ipv6 router ospf area 0.0.0.0 tag 100`	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#interface eth2`	Enter interface mode.
`R1(config-if)#ipv6 address 6111::1/64`	Configure the IPv6 address of the interface.
`R1(config-if)#ipv6 router ospf area 0.0.0.0 tag 100`	Enable OSPFv3 routing on an interface, and assign the Area ID 0.

© 2026 IP Infusion Inc. Proprietary

`R1(config-if)# ipv6 ospf cost 2`	Configuring ospf cost as 2
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#interface eth3`	Enter interface mode.
`R1(config-if)#ipv6 address 1311::1/64`	Configure the IPv6 address of the interface.
`R1(config-if)#ipv6 router ospf area 0.0.0.0 tag 100`	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
`R1(config-if)# ipv6 ospf cost 3`	Configuring ospf cost as 3
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#interface eth4`	Enter interface mode.
`R1(config-if)#ipv6 router ospf area 0.0.0.0 tag 100`	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
`R1(config-if)#ipv6 address 101::2/64`	Configure the IPv6 address of the interface.
`R1(config-if)#ipv6 policy route-map 124`	Attach PBR on the ingress interface
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-if)#exit`	Exit interface mode.

Validation

```
R1#show track
TRACK Id: 8
  IP SLA 8 reachability
  Reachability is UP
    7 changes, last change : 2025 Jun 19 09:05:07
Track interface :

R1#show track summary
Object Tracking Summary
ID       Type        Type-Identifier       State
-----------------------------------------------------
8      ip-sla            8               UP

R1#show nsm obj-track-table
Object Track Table
Track Id Track Status Last Updated
8        Up           2025 Jun 19 09:09:07


R1#show hsl pbr

PBR name : 123, FIB ID : 0, Family : IP
Rule count : 1, Interface count : 1

Interface(RuleCnt) : eth4(2)

PBR rule number: 10
  MATCH details:
    IP-LIST   - ACL-name : 124 (IfRefCnt:1)
  SET details:
    Action : Route
    Next-hop count : 1
      Next-hop 1311::2 (eth1), VALID, INSTALLED, NH 0x1c19fbe8, TRACKED
    Load balance : NO
    Egress object ID : 0x2000ccdd
```

TOS based Queue Distribution Configuration

This feature allows both type of configuration (precedence or DSCP). Before the implementation of this feature only allows DSCP profiles.The most significant three bits of the DSCP are Class Selector bits, which provides backward compatibility with ip precedence.

We will utilize these DSCP bits to set Precedence. For example when user configure Precedence 101 in CLI, we just need to convert this value to DSCPs(101000 ~ 101111), then normally set the DSCP.

Indicate the standards to be imposed or other all-embracing blanket requirements.

On the basis of already existing DSCP profiles, new precedence profiles are introducing as qos map-profile in this section.

The working on these precedence profile is similar to that of DSCP but now we can permit both the configurations according to our requirement.

There are 2 Qos-map Precedence profiles introduced:

- Precedence-to-queue
- Precedence-to-precedence

Precedence-to-queue

This is an ingress profile similar to dscp-to-queue. This profile will map precedence value 0-7 to queue 0-7.

Topology

Figure 17. Precedence-to-queue topology

Configuration

Basic configuration

`#configure terminal`	Enter configure mode
`(config)#enable qos`	Enable Qos
`(config)#exit`	Exit configure mode
`(config)# qos profile precedence-to-queue Prec1_que1`	Configure qos-map profile precedence-to-queue with Profile name as Prece1_que1
`(config-ingress-prec-map)# precedence 1 queue 2`	Configure mapping of prece 1 to queue 2, so that traffic with precedence 1 should take queue as q2
`(config)#exit`	Exit configure mode
`(config)# interface xe1`	Configure ingress interface xe1
`(config-if)#qos map-profile precedence-to-queue Prec1_que1`	Map the profile Prec1_que1 to interface xe1 so that ingress traffic hitting the xe1 interface with prece as 1 should take queue q2

© 2026 IP Infusion Inc. Proprietary

(config-if)#commit	Commit the configuration
(config-if)#end	Exit the configuration.

Validation

```
#show run int xe1
!
interface xe1
 qos map-profile precedence-to-queue Prec1_que1
!

#show qos-profile type precedence-to-queue
profile name: Prec1_que1
profile type: precedence-to-queue
profile attached to 1 instances
configured mapping:
 precedence 1 queue 2
 Detailed mapping:
 -------+------------------
  INPUT |    OUTPUT
 -------+------------------
  PREC  | Queue | Out PREC
 -------+-------+----------
  0       0       0
  1       1       2
  2       2       0
  3       3       0
  4       4       0
  5       5       0
  6       6       0
  7       7       0


R1#show run qos
qos enable
!
qos profile precedence-to-queue Prec1_que1
 precedence 1 queue 2
!
interface xe1
 qos map-profile precedence-to-queue Prec1_que1
!
```

Precedence-to-precedence

This is an egress profile similar to dscp-to-dscp. This profile will map precedence value 0-7 to precedence value 0-7. It will change the egress queue according to new mapped precedence value.

Topology

Figure 18. Precedence-to-precedence topology

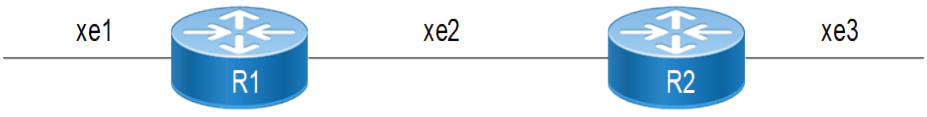

Basic configuration

`#configure terminal`	Enter configure mode
`(config)#enable qos`	Enable Qos
`(config)#exit`	Exit configure mode
`(config)# qos profile precedence-to-precedence prec_prec2`	Configure qos-map profile precedence-to-precedence with Profile name as prec_prec2
`(config-ingress-prec-map)# precedence 2 precedence 5`	Configure mapping of prece 2 to precedence 5, so that traffic with precedence 2 should remark to precedence 5 and should take corresponding queue q5 on next router
`(config)#exit`	Exit configure mode
`(config)# interface xe2`	Configure ingress interface xe2
`(config-if)#qos map-profile precedence-to-queue prec_prec2`	Map the profile prec_prec2 to egress interface xe2 of 1st node.
`(config-if)#commit`	Commit the configuration
`(config-if)#end`	Exit the configuration.

Validation

```
#show run qos
qos enable
!
qos profile precedence-to-precedence prec_prec2
 precedence 2 precedence 5
!
!
!
!
interface xe2
 qos map-profile precedence-to-precedence prec_prec2
!

#show qos-profile prec_prec2
profile name: prec_prec2
profile type: precedence-to-precedence
profile attached to 1 instances
configured mapping:
 precedence 2 precedence 5
Detailed mapping:
 -------+--------
  INPUT | OUTPUT
 -------+--------
  PREC  | PREC
 -------+--------
  0       0
  1       1
  2       5
  3       3
  4       4
  5       5
  6       6
  7       7
```

Internet Protocol SLA Configuration

The Internet Protocol Service Level Agreement (IP SLA) is an active method for monitoring and reliably reporting network performance. The term "active" refers to IP SLA's ability to generate and actively monitor traffic across the network on a continuous basis. An IP SLA Router is capable of generating traffic and providing real-time reporting.

IP SLA configuration consists of two components: the IP SLA router, which generates traffic, and the IP SLA Responder, which can be any device, not limited to just routers. While the IP SLA Responder is not mandatory for IP SLA functionality, it enhances the ability to gather more detailed information and generate more comprehensive reports.

Once IP SLA operations are configured, scheduling the operation is required to start capturing statistics and collecting error data. The operation can be scheduled to start immediately or at a specified month, day, and hour. A "pending" option is also available to schedule the operation for a later time. This "pending" state is visible through SNMP and is utilized when an operation is designed to be triggered by a certain threshold.

You can schedule a single IP SLA operation or a group of operations to start at once.

Note: IP SLA sessions on Edgecore AS7316-26XB switches are limited to 500 sessions. This limit may vary depending on the device's capacity and performance.

Topology

Figure 19. 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.

© 2026 IP Infusion Inc. Proprietary

`(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
```

EVPN AF Route-Maps and Route Filtering

Overview

The EVPN AF Route-Maps and Route Filtering feature extends the existing route-map framework to the L2VPN EVPN address family in BGP. It allows users to filter, modify, or manage EVPN routes exchanged between BGP peers.

A route-map defines the criteria to match specific route attributes and actions to set or modify those attributes. Route-maps can be applied in the incoming (IN) or outgoing (OUT) direction on a BGP neighbor or peer group.

- OUT direction: Processes EVPN routes before BGP advertises them.
- IN direction: Processes routes after BGP receives them.

The framework introduces EVPN-specific match conditions, such as route type and MAC lists, enabling precise control over route propagation and reducing control-plane overhead in large BGP-EVPN networks.

Feature Characteristics

- Supports route-map configuration under the L2VPN EVPN address family.
- Applies route-maps in both IN and OUT directions.
- Integrates with the existing route-map framework with EVPN extensions.
- Supports numbered and unnumbered L2VPN modes.
- Evaluates routes using match and set conditions.
- Adds EVPN-specific match capabilities:
  - match evpn-route-type : Type-1 to Type-5, MAC-only, MAC-IP routes.
- Match mac-list : Permit or deny specific MAC addresses.
- Supports standard BGP route-map matches like ASpath, next-hop, route-target and other match options.
- Controls route propagation direction and scope.
- Reduces unnecessary route advertisements across spines, leaves, and superspines.

Benefits

- Provides granular EVPN route control.
- Allows selective advertisement, acceptance, or rejection of routes.
- Enhances scalability and performance by reducing redundant updates.
- Simplifies multi-tenant isolation via communities or route-targets.
- Ensures consistent configuration across IPv4, IPv6, and EVPN AFs.
- Reduces control-plane load and processing time for operational efficiency.

Prerequisites

Before configuring EVPN AF Route-Maps and Route Filtering, ensure the following conditions are met:

© 2026 IP Infusion Inc. Proprietary

- BGP EVPN Setup:
  - BGP sessions between all relevant PEs, RRs, and leaf/spine nodes must be established.
  - The L2VPN EVPN address family must be enabled on all participating BGP neighbors.

Limitation

- The set option is not supported for import-map in VRF configurations.
- The AIGP path attribute is not supported for EVPN address-family routes. Therefore, when configuring a neighbor route map under the EVPN address family or a route map used in a VRF export map, the **`set aigp-metric <>`** command is not applicable to EVPN routes.
- Under route-map configuration, **`set ip nexthop`** is not supported for EVPN address family routes.

Configuration

This section describes the configuration of EVPN Address Family (AF) with route filtering using route-maps on PE routers.

Route-maps are applied in the EVPN AF to control route import or export policies and to modify all the supported attributes.

Topology

The topology illustrates an EVPN network using MPLS underlay and BGP control plane for end-to-end connectivity between CE1 and CE2.

Figure 20. EVPN AF Route Map

PE1 and PE3 function as Provider Edge routers participating in EVPN BGP sessions, while RR acts as the Route Reflector to distribute EVPN routes between them.

PE1–RR uses LDP with OSPF, and RR–PE3 uses RSVP with ISIS for MPLS transport. BGP Labeled Unicast (BGP-LU) establishes end-to-end label distribution between PEs.

Route-maps are applied for route control:

- PE1 applies an OUT route-map for route advertisement.
- RR applies an IN route-map for route reflection.
- PE3 applies an IN route-map for route import and CE2 advertisement.

Traffic from CE1 to CE2 traverses the MPLS core, with BGP EVPN managing MAC and IP route exchange between PEs.

Router or Node Configuration Steps

Perform the following configuration steps to apply route-maps under the BGP EVPN address family, enabling selective advertisement and acceptance of EVPN routes between peers.

Note: Before configuration meet all Prerequisites (page 261).

1. Configure MAC Access List on PE1.
   Define a MAC access list to restrict specific MAC addresses from being advertised through EVPN.

```
mac-list mac_list1
 seq 10 deny 6821.5f1f.5220 ffff.ffff.ffff
 seq 20 permit 0000.0000.0000 0000.0000.0000
```

2. Configure Outbound Route-Map on PE1.
   a. Define a route-map named **`rm_evpn_out`** to control the EVPN route advertisement.
   b. Match different EVPN route types and set appropriate BGP attributes for each type.

```
route-map rm_evpn_out permit 10
 match evpn-route-type type-1
 set local-preference 200
 set community 65000:100 additive

route-map rm_evpn_out permit 20
 match evpn-route-type type-2
 match mac address list mac_list1
 set local-preference 150
 set metric 50
 set community 65000:200 additive

route-map rm_evpn_out permit 30
 match evpn-route-type type-3
 set local-preference 100
 set origin igp
 set community 65000:300 additive

route-map rm_evpn_out permit 1000
```

   c. Apply the route-map to the BGP EVPN address family for outbound updates to the Route Reflector.

```
router bgp 65000
 address-family l2vpn evpn
  neighbor 3.3.3.3 activate
  neighbor 3.3.3.3 route-map rm_evpn_out out
 exit-address-family
```

3. Configure Inbound Route-Map on PE3.
   a. Create a route-map named **`rm_evpn_in`** to control inbound EVPN route processing.
   b. Match specific EVPN route types and adjust route selection parameters.

```
route-map rm_evpn_in permit 10
 match evpn-route-type type-1
 set weight 200
```

   c. Apply the route-map to the BGP EVPN address family for inbound routes from the Route Reflector.

```
router bgp 65000
 address-family l2vpn evpn
  neighbor 3.3.3.3 activate
  neighbor 3.3.3.3 route-map rm_evpn_in in
 exit-address-family
```

4. Configure the RR for EVPN.
   a. Enable EVPN under the BGP configuration.
   b. Configure both PE1 and PE3 as route-reflector clients.
   c. Activate the EVPN address family for both peers.

```
router bgp 65000
 bgp router-id 3.3.3.3
 address-family l2vpn evpn
```

© 2026 IP Infusion Inc. Proprietary

```
   neighbor 1.1.1.1 activate
   neighbor 1.1.1.1 route-reflector-client
   neighbor 5.5.5.5 activate
   neighbor 5.5.5.5 route-reflector-client
  exit-address-family
```

Validation

- Verify BGP Neighbor Status.

```
PE1#show ip bgp labeled-unicast summary
BGP router identifier 1.1.1.1, local AS number 65000
BGP table version is 3
1 BGP AS-PATH entries
3 BGP community entries

Neighbor           V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd
Desc
3.3.3.3            4 65000       1312       1302       3      0      0  04:59:49             2

Total number of neighbors 1

Total number of Established sessions 1
```

- Verify that all BGP EVPN neighbor sessions are established.

```
show bgp l2vpn evpn summary
BGP router identifier 1.1.1.1, local AS number 65000
BGP table version is 15
1 BGP AS-PATH entries
3 BGP community entries

Neighbor          V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd
 AD  MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
3.3.3.3            4
65000       1312       1302      15      0      0  04:59:49            5      1      3      1
   0      0

Total number of neighbors 1

Total number of Established sessions 1
```

- Verify that the route-maps are correctly applied.

```
PE1#show route-map
route-map rm_lu1, permit, sequence 10
  Match clauses:
    ip address prefix-list: lu_pref1
  Set clauses:
    local-preference 300
route-map rm_lu1, deny, sequence 11
  Match clauses:
    ip address prefix-list: lu_pref2
  Set clauses:
    originator-id 15.15.15.15
route-map rm_evpn_out, permit, sequence 10
  Match clauses:
    evpn-route-type: type-1
  Set clauses:
    local-preference 200
    community 65000:100  (additive)
route-map rm_evpn_out, permit, sequence 20
  Match clauses:
    evpn-route-type: type-2
    mac address list: mac_list1
  Set clauses:
    local-preference 150
    metric 50
```

© 2026 IP Infusion Inc. Proprietary

```
    community 65000:200  (additive)
route-map rm_evpn_out, permit, sequence 30
  Match clauses:
    evpn-route-type: type-3
  Set clauses:
    local-preference 100
    origin igp
    community 65000:300  (additive)
route-map rm_evpn_out, permit, sequence 1000
  Match clauses:
  Set clauses:
```

- Verify EVPN-MPLS tunnel status.

```
PE1#show evpn mpls xconnect tunnel
EVPN-MPLS Network tunnel Entries
Source           Destination      Status        Up/Down       Update        local-evpn-id remote-
evpn-id Ext-Color FAT
================================================================================================
====================
1.1.1.1          5.5.5.5          Installed     04:59:49      04:59:49      501           133501
          ---

Total number of entries are 1

PE1#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source           Destination      Status        Up/Down       Update        evpn-id       Local-
Leaf  Remote-Leaf  Ext-Color  FAT
================================================================================================
================================
1.1.1.1          5.5.5.5          Installed     04:59:49      04:59:49      205             ---
         ---           ---

Total number of entries are 1
```

- Verify EVPN MPLS Interfaces.

```
PE1#show evpn mpls
EVPN-MPLS Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VPN-ID   EVI-Name     EVI-Type Type Interface ESI                          VLAN      DF-Status
Src-Addr         Dst-Addr
________________________________________________________________________________________________
_________________________________
205      ----         L2       NW   ----      ----                         ----      ----
      1.1.1.1          5.5.5.5          --------------
205      ----         --       AC   xe25.205  --- Single Homed Port ---    ----      ----      -
---             ----

Total number of entries are 2

Note: Refer sub-interface config for VLAN information.
```

- Verify BGP EVPN routes.

```
PE1#show bgp l2vpn evpn
BGP table version is 15, local router ID is 1.1.1.1
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
```

```
Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next
Hop            Metric    LocPrf       Weight     Path  Peer          Encap

RD[1.1.1.1:105] VRF[vrf105]
*>    [2]:[0]:[205]:[48,0011:2233:4455]:[32,98.98.101.1]:[17]
                      1.1.1.1              0        100       32768  i        -  ----------
      MPLS
*>    [2]:[0]:[205]:[48,0011:2233:4455]:[32,104.104.103.1]:[17]
                      1.1.1.1              0        100       32768  i        -  ----------
      MPLS
* i l [2]:[0]:[205]:[48,0011:2233:4567]:[32,99.99.101.1]:[17]
                      5.5.5.5              0        100       0    i        -
  3.3.3.3         MPLS
* i l [2]:[0]:[205]:[48,0011:2233:4567]:[32,104.104.104.1]:[17]
                      5.5.5.5              0        100       0    i        -
  3.3.3.3         MPLS
*>    [2]:[0]:[205]:[48,6821:5f1f:5220]:[32,103.103.102.1]:[17]
                      1.1.1.1              0        100       32768  i        -  ----------
      MPLS
* i l [2]:[0]:[205]:[48,e8c5:7a69:45ed]:[32,103.103.103.1]:[17]
                      5.5.5.5              0        100       0    i        -
  3.3.3.3         MPLS
*>    [3]:[205]:[32,1.1.1.1]
                      1.1.1.1              0        100       32768  i        -  ----------
      MPLS
* i   [3]:[205]:[32,5.5.5.5]
                      5.5.5.5              0        100       0    i        -
  3.3.3.3         MPLS

RD[1.1.1.1:501] VRF[Eline501]
*>    [1]:[0]:[501]:[18]
                      1.1.1.1              0        100       32768  i        -  ----------
      MPLS
* i l [1]:[0]:[133501]:[18]
                      5.5.5.5              0        100       0    i        -
  3.3.3.3         MPLS

RD[5.5.5.5:105]
*>i l [2]:[0]:[205]:[48,0011:2233:4567]:[32,99.99.101.1]:[17]
                      5.5.5.5              0        100       0    i        -
  3.3.3.3         MPLS
*>i l [2]:[0]:[205]:[48,0011:2233:4567]:[32,104.104.104.1]:[17]
                      5.5.5.5              0        100       0    i        -
  3.3.3.3         MPLS
*>i l [2]:[0]:[205]:[48,e8c5:7a69:45ed]:[32,103.103.103.1]:[17]
                      5.5.5.5              0        100       0    i        -
  3.3.3.3         MPLS
*>i   [3]:[205]:[32,5.5.5.5]
                      5.5.5.5              0        100       0    i        -
  3.3.3.3         MPLS

RD[5.5.5.5:501]
*>i l [1]:[0]:[133501]:[18]
                      5.5.5.5              0        100       0    i        -
  3.3.3.3         MPLS

Total number of prefixes 15
```

- Verify MAC/IP table.

```
PE1#show bgp l2vpn evpn mac-ip
BGP table version is 15, local router ID is 1.1.1.1
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
Description : Ext-Color - Extended community color


RD[1.1.1.1:105] VRF[vrf105]:
     ESI                              Eth-Tag     Mac-Address    IP-
Address      VNID/LABEL     L3VNID    Nexthop         GW-Type         Peer            Encap
*>   0                                205         0011:2233:4455
98.98.101.1     17             0         1.1.1.1         --                ----------      MPLS
*>   0                                205         0011:2233:4455
104.104.103.1   17             0         1.1.1.1         --                ----------      MPLS
* il 0                                205         0011:2233:4567
99.99.101.1     17             0         5.5.5.5         --                3.3.3.3         MPLS
* il 0                                205         0011:2233:4567
104.104.104.1   17             0         5.5.5.5         --                3.3.3.3         MPLS
*>   0                                205         6821:5f1f:5220
103.103.102.1   17             0         1.1.1.1         --                ----------      MPLS
* il 0                                205         e8c5:7a69:45ed
103.103.103.1   17             0         5.5.5.5         --                3.3.3.3         MPLS

RD[5.5.5.5:105]
     ESI                              Eth-Tag     Mac-Address    IP-
Address      VNID/LABEL     L3VNID    Nexthop         GW-Type         Peer            Encap
*>il 0                                205         0011:2233:4567
99.99.101.1     17             0         5.5.5.5         --                3.3.3.3         MPLS
*>il 0                                205         0011:2233:4567
104.104.104.1   17             0         5.5.5.5         --                3.3.3.3         MPLS
*>il 0                                205         e8c5:7a69:45ed
103.103.103.1   17             0         5.5.5.5         --                3.3.3.3         MPLS

PE1#show evpn mpls mac-table
================================================================================================================
================================================================
                                                    EVPN MPLS MAC Entries
================================================================================================================
================================================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-
Ip/ESI                                   Type            Status     MAC move AccessPortDesc
LeafFlag
________________________________________________________________________________________________________________
_____________________________________________________________

205        irb305    ----      ----      0011.2233.4455
1.1.1.1                                       Static Local    -------    0        -------
          ----
205        ----      ----      ----      0011.2233.4567
5.5.5.5                                       Static Remote   -------    0        -------
          ----
205        irb305    ----      ----      6821.5f1f.5220
1.1.1.1                                       Static Local    -------    0        -------
          ----
205        ----      ----      ----      e8c5.7a69.45ed
5.5.5.5                                       Static Remote   -------    0        -------
          ----

Total number of entries are : 4

PE1#show evpn mpls arp-cache
MPLS-EVPN ARP-CACHE Information
================================
EVPN-ID   Ip-Addr          Mac-Addr        Type        Age-Out   Retries-Left
________________________________________________________________________________
205      98.98.101.1      0011.2233.4455 Static Local   ----
205      99.99.101.1      0011.2233.4567 Static Remote  ----
205      103.103.102.1    6821.5f1f.5220 Static Local   ----
```

```
205      103.103.103.1    e8c5.7a69.45ed Static Remote  ----
205      104.104.103.1    0011.2233.4455 Static Local   ----
205      104.104.104.1    0011.2233.4567 Static Remote  ----
Total number of entries are 6

PE1#show bgp l2vpn evpn detail
BGP route entry for prefix : [2]:[0]:[205]:[48,0011:2233:4455]:[32,98.98.101.1]:[17]
Route-Distinguisher : [1.1.1.1:105] VRF : vrf105
Flags : Valid, Selected
Nexthop : 1.1.1.1 MED value: 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
Weight : 32768, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 13:46:13 2025
Peer : Local
BGP route entry for prefix : [2]:[0]:[205]:[48,0011:2233:4455]:[32,104.104.103.1]:[17]
Route-Distinguisher : [1.1.1.1:105] VRF : vrf105
Flags : Valid, Selected
Nexthop : 1.1.1.1 MED value: 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
Weight : 32768, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 13:46:13 2025
Peer : Local
BGP route entry for prefix : [2]:[0]:[205]:[48,0011:2233:4567]:[32,99.99.101.1]:[17]
Route-Distinguisher : [1.1.1.1:105] VRF : vrf105
Flags : Valid, IBGP, Labelled, Labelled
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
BGP route entry for prefix : [2]:[0]:[205]:[48,0011:2233:4567]:[32,104.104.104.1]:[17]
Route-Distinguisher : [1.1.1.1:105] VRF : vrf105
Flags : Valid, IBGP, Labelled, Labelled
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
Weight : 0, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
BGP route entry for prefix : [2]:[0]:[205]:[48,6821:5f1f:5220]:[32,103.103.102.1]:[17]
Route-Distinguisher : [1.1.1.1:105] VRF : vrf105
Flags : Valid, Selected
Nexthop : 1.1.1.1 MED value: 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
Weight : 32768, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 13:46:13 2025
Peer : Local
BGP route entry for prefix : [2]:[0]:[205]:[48,e8c5:7a69:45ed]:[32,103.103.103.1]:[17]
Route-Distinguisher : [1.1.1.1:105] VRF : vrf105
Flags : Valid, IBGP, Labelled, Labelled
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
Weight : 0, Local Preference :100
AS Path : Local
```

```
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
BGP route entry for prefix : [3]:[205]:[32,1.1.1.1]
Route-Distinguisher : [1.1.1.1:105] VRF : vrf105
Flags : Valid, Selected
Nexthop : 1.1.1.1 MED value: 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS
Weight : 32768, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 13:46:13 2025
Peer : Local
BGP route entry for prefix : [3]:[205]:[32,5.5.5.5]
Route-Distinguisher : [1.1.1.1:105] VRF : vrf105
Flags : Valid, IBGP
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS
Weight : 0, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
BGP route entry for prefix : [1]:[0]:[501]:[18]
Route-Distinguisher : [1.1.1.1:501] VRF : Eline501

Flags : Valid, Selected
Nexthop : 1.1.1.1 MED value: 0
Community :
Extended Community: RT:501:501 Encapsulation:MPLS ESI-Label:0 Control-Flags,MTU:2,9216
Weight : 32768, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 15:57:57 2025
Peer : Local
BGP route entry for prefix : [1]:[0]:[133501]:[18]
Route-Distinguisher : [1.1.1.1:501] VRF : Eline501
Flags : Valid, IBGP, Labelled, Labelled
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:501:501 Encapsulation:MPLS ESI-Label:0 Control-Flags,MTU:2,9216
Weight : 0, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
BGP route entry for prefix : [2]:[0]:[205]:[48,0011:2233:4567]:[32,99.99.101.1]:[17]
Route-Distinguisher : [5.5.5.5:105]
Flags : Valid, Selected, IBGP, Labelled, Labelled
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
Weight : 0, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
BGP route entry for prefix : [2]:[0]:[205]:[48,0011:2233:4567]:[32,104.104.104.1]:[17]
Route-Distinguisher : [5.5.5.5:105]
Flags : Valid, Selected, IBGP, Labelled, Labelled
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
Weight : 0, Local Preference :100
AS Path : Local
Origin : IGP
```

```
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
BGP route entry for prefix : [2]:[0]:[205]:[48,e8c5:7a69:45ed]:[32,103.103.103.1]:[17]
Route-Distinguisher : [5.5.5.5:105]
Flags : Valid, Selected, IBGP, Labelled, Labelled
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
Weight : 0, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
BGP route entry for prefix : [3]:[205]:[32,5.5.5.5]
Route-Distinguisher : [5.5.5.5:105]
Flags : Valid, Selected, IBGP
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS
Weight : 0, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025

Peer : 3.3.3.3
BGP route entry for prefix : [1]:[0]:[133501]:[18]
Route-Distinguisher : [5.5.5.5:501]
Flags : Valid, Selected, IBGP, Labelled, Labelled
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:501:501 Encapsulation:MPLS ESI-Label:0 Control-Flags,MTU:2,9216
Weight : 0, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
Total number of prefixes 15
PE1#show bgp l2vpn evpn vrf vrf105
BGP table version is 1, local router ID is 1.1.1.1
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
Description : Ext-Color - Extended community color
[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route
Network Next
Hop Metric LocPrf Weight Path Peer Encap
*> [2]:[0]:[205]:[48,0011:2233:4455]:[32,98.98.101.1]:[17]
1.1.1.1 0 100 32768 i - --------
-- MPLS
*> [2]:[0]:[205]:[48,0011:2233:4455]:[32,104.104.103.1]:[17]
1.1.1.1 0 100 32768 i - --------
-- MPLS
* i l [2]:[0]:[205]:[48,0011:2233:4567]:[32,99.99.101.1]:[17]
5.5.5.5 0 100 0 i -
3.3.3.3 MPLS
* i l [2]:[0]:[205]:[48,0011:2233:4567]:[32,104.104.104.1]:[17]
5.5.5.5 0 100 0 i -
3.3.3.3 MPLS
*> [2]:[0]:[205]:[48,6821:5f1f:5220]:[32,103.103.102.1]:[17]
1.1.1.1 0 100 32768 i - --------
-- MPLS
* i l [2]:[0]:[205]:[48,e8c5:7a69:45ed]:[32,103.103.103.1]:[17]
```

```
5.5.5.5 0 100 0 i -
3.3.3.3 MPLS
*> [3]:[205]:[32,1.1.1.1]
1.1.1.1 0 100 32768 i - --------
-- MPLS
* i [3]:[205]:[32,5.5.5.5]
5.5.5.5 0 100 0 i -
3.3.3.3 MPLS
Total number of prefixes 8
```

Implementation Examples

- EVPN Route Control Between Leaf and Spine Nodes : Leaf nodes may advertise routes that other spines do not require. You can apply a route-map to filter those routes based on AS numbers or route types.
- Community-Based Route Filtering: When VRF routes are redistributed into EVPN, you can tag them with specific communities and use route-maps to advertise only selected routes.

Commands

The EVPN AF route-maps and route filtering feature introduces the following commands:

mac-list

Use this command to create a new MAC list that specifies MAC addresses with an associated mask and defines an action (permit or deny).

Use the **no** parameter of this command to unconfigure.

Command Syntax

Mac List:

```
mac-list WORD

no mac-list WORD
```

Mac list entry:

```
permit/deny XXXX.XXXX.XXXX XXXX.XXXX.XXXX

no permit/deny XXXX.XXXX.XXXX XXXX.XXXX.XXXX
```

Parameters

mac-list
Build a MAC list

WORD
Name of MAC list

Permit
Permits the MAC Address (XXXX.XXXX.XXXX) in HHHH.HHHH.HHHH format

Deny
Denies the MAC Address (XXXX.XXXX.XXXX) in HHHH.HHHH.HHHH format

Default

None

Command Mode

Configure mode and Mac-list mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example illustrates how to specifies MAC addresses with an associated mask and defines an action (permit or deny):

```
ocnos#configure terminal
ocnos(config)#mac-list mlist1
ocnos(config-mac-list)permit 1111.1111.3322 1100.0011.0000
```

© 2026 IP Infusion Inc. Proprietary

match mac address list

Use this command to match a mac address list.

Use the **no** parameter of this command to turn off the matching.

Command Syntax

```
match mac address list WORD

no match mac address list WORD
```

Parameters

WORD

Defines the MAC list name

Default

None

Command Mode

Router Map mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example illustrates how to match a mac address list:

```
R1#configure terminal
R1(config)#route-map rmap1 permit 10
R1(config-route-map)# match mac address list mlist1
```

© 2026 IP Infusion Inc. Proprietary

match evpn-route-type

Use this command to match an EVPN route by its route type. The supported route types are: **`type-1`**, **`type-2`**, **`type-3`**, **`type-4`**, **`type-5`**, **`type-2-MAC-IP`**, and **`type-2-MAC-ONLY`**.

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

Command Syntax

```
match evpn-route-type (type-1|type-2|type-3|type-4|type-5|type-2-MAC-IP|type-2-MAC-ONLY|)

no match evpn-route-type (type-1|type-2|type-3|type-4|type-5|type-2-MAC-IP|type-2-MAC-ONLY|)
```

Parameters

type-1
Matches Ethernet Auto-Discovery (AD) routes

type-2
Matches both MAC-only and MAC/IP advertisement routes

type-2-MAC-IP
Matches MAC/IP advertisement routes

type-2-MAC-ONLY
Matches MAC-only advertisement routes

type-3
Matches Multicast Ethernet Tag (MET) routes

type-4
Matches Ethernet Segment Identifier (ESI) routes

type-5
Matches IP Prefix routes

Default

None

Command Mode

Router Map mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example illustrates how to match an EVPN route by its route type:

```
R1#configure terminal
R1(config)#route-map rmap1 permit 10
R1(config-route-map)# match evpn-route-type type-2
```

© 2026 IP Infusion Inc. Proprietary

show mac-list

Use this command to view the detailed information about a specific mac-list.

Command Syntax

```
show mac-list detail WORD
```

Parameters

WORD

Specifies the name of the MAC list for which detailed information is to be displayed.

Default

None

Command Mode

Privileged execution mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example illustrates how to view the detailed information about a specific mac-list:

```
ocnos#show mac-list detail mlist1
mac-list mlist1:
   count: 2, sequences: 3 - 5
    ripd:
    ripngd:
    ospfd:
    ospf6d:
    ldpd:
    bgpd:
     seq 3 permit  0000.aaaa.bbbb ffff.ffff.ffff (hit count: 10, refcount: 21)
     seq 5 deny  0000.ccaa.bbbb ffff.ffff.ffff (hit count: 1, refcount: 2)
```

Glossary

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

Key Terms/Acronym	Description
Address Family Identifier / Subsequent AFI (AFI/SAFI)	Specifies network layer protocol and route type.
Border Gateway Protocol (BGP)	Routing protocol used to exchange routing information across autonomous systems.
Ethernet VPN (EVPN)	BGP-based control plane for Ethernet multipoint services.

© 2026 IP Infusion Inc. Proprietary

FUNDAMENTAL LAYER 3 COMMAND REFERENCE

Fundamental Layer 3 Commands 279
automatic-router-id-selection enable 280
clear ip route 281
clear ip route kernel 282
clear ip route vrf 283
clear router-id 284
debug rib 285
description 287
fib retain 288
ip route 290
ip route vrf 293
ip urpf allow-default 294
ip urpf enable 295
ip verify unicast source reachable-via 296
ip vrf 297
ipv6 route 298
max-static-routes 300
maximum-paths 301
router-id 302
show debugging rib 303
show ip route track-table 304
show ip rpf 305
show ipv6 route track-table 307
show ipv6 rpf 308
show rib txlist 310
show router-id 312
show running-config router 313
show running-config router-id 314
show running-config urpf 315
show running-config vrf 316
snmp restart rib 317
IP Service Level Agreements Commands 318
clear ip sla statistics 319
frequency 320
icmp-echo 321

© 2026 IP Infusion Inc. Proprietary

ip sla 322
ip sla schedule 323
show ip sla statistics 324
show ip sla summary 326
show running-config ip sla 327
threshold 328
timeout 329
Object Tracking Commands 330
delay up down 331
object-tracking 332
show running-config track 334
show track 335
show track <1-500> 336
show track summary 337
track ip sla reachability 338
Route-Map Commands 339
continue 341
match as-path 343
match community 344
match extcommunity 345
match interface 346
match ip address 347
match ip address prefix-list 348
match ip next-hop 349
match ip next-hop prefix-list 350
match ip peer 351
match ipv6 address 352
match ipv6 address prefix-list 353
match ipv6 next-hop 354
match ipv6 next-hop prefix-list 355
match ipv6 peer 356
match metric 357
match origin 358
match route-map 359
match route-type 361
match tag 362
route-map 363
set aggregator 364
set as-path 365
set atomic-aggregate 366

© 2026 IP Infusion Inc. Proprietary

set comm-list 367
set community 368
set dampening 370
set extcommunity 371
set interface null0 373
set ip next-hop 374
set ipv6 next-hop 376
set ip next-hop self 377
set large-community 378
set level 380
set local-preference 381
set metric 382
set metric-type 383
set origin 384
set originator-id 385
set tag 386
set vpnv4 next-hop 387
set weight 388
show route-map 389
show running-config route-map 390
Source Interface Commands 391
ip source-interface 392
ipv6 source-interface 394
show ip source-interface detail 396
show ipv6 source-interface detail 398
show running-config ip source-interface 400
show running-config ipv6 source-interface 401

© 2026 IP Infusion Inc. Proprietary

Fundamental Layer 3 Commands

This section describes the fundamental Layer 3 commands:

automatic-router-id-selection enable 280
clear ip route 281
clear ip route kernel 282
clear ip route vrf 283
clear router-id 284
debug rib 285
description 287
fib retain 288
ip route 290
ip route vrf 293
ip urpf allow-default 294
ip urpf enable 295
ip verify unicast source reachable-via 296
ip vrf 297
ipv6 route 298
max-static-routes 300
maximum-paths 301
router-id 302
show debugging rib 303
show ip route track-table 304
show ip rpf 305
show ipv6 route track-table 307
show ipv6 rpf 308
show rib txlist 310
show router-id 312
show running-config router 313
show running-config router-id 314
show running-config urpf 315
show running-config vrf 316
snmp restart rib 317

© 2026 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)#
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

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

© 2026 IP Infusion Inc. Proprietary

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

Execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear ip route kernel
```

© 2026 IP Infusion Inc. Proprietary

clear ip route vrf

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

Execution 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`.

Note: Router-ID is non-preventive and changes at the routing protocol level or removal/modification of the address used for router-ID selection are not automatically reflected. The system retains the previously selected router-ID, use the `clear router-id` command to force re-selection.

Command Syntax

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

Parameters

VRFNAME
VPN routing/forwarding instance name.

Command Mode

Execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

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

© 2026 IP Infusion Inc. Proprietary

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

© 2026 IP Infusion Inc. Proprietary

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 execution mode

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

Virtual-router instance 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
```

© 2026 IP Infusion Inc. Proprietary

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.

Below table 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 (page 282) 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.

© 2026 IP Infusion Inc. Proprietary

Examples

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

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

ip route vrf

Use this command to configure a tracking-id to ipv4 static route in non-default vrf.

Use the `no` form of this command to delete the tracking id

Command Syntax

```
ip route vrf <vrf-name> A.B.C.D/M (A.B.C.D|IFNAME){<track <1-500>}
```

Parameters

vrf-name
Vrf name

A.B.C.D/M
Destination prefix with subnet in ipv4 format

A.B.C.D
Gateway address in ipv4 format

IFNAME
Gateway address in interface name format

<1-500>
Identifier for the tracked object

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 5.1.

Example

```
#configure terminal
(config)# ip route vrf vrf1 40.1.1.0/24 eth1 track 10
(config-object-track)#commit

(config)#no ip route vrf vrf1 40.1.1.0/24 eth1 track 10
(config)#commit
```

© 2026 IP Infusion Inc. Proprietary

ip urpf allow-default

Use this command to enable default route check for uRPF mode.

Use `no` form of this command to disable default route check for uRPF mode.

Command Syntax

This command applies only to Qumran platforms.

```
ip urpf allow-default
no ip urpf allow-default
```

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 allow-default

(config)#no ip urpf allow-default
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 | 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

Default

None

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 rx
(config-if)#no ip verify unicast source reachable-via
```

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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. Added a new parameter `recursive` in OcNOS version 6.6.0.

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

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)#
```

© 2026 IP Infusion Inc. Proprietary

show debugging rib

Use this command to display debug settings.

Command Syntax

```
show debugging rib
```

Parameters

None

Command Mode

Privileged execution mode and Execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
# show debugging rib
```

© 2026 IP Infusion Inc. Proprietary

show ip route track-table

Use this command to display ipv4 static routes that are associated to a tracking id across all vrfs.

Command Syntax

```
show ip route track-table
```

Parameters

None

Default

None

Command Mode

Privileged execution mode and Execution mode

Applicability

This command is introduced in OcNOS version 5.1.

Examples

```
#show ip rou track-table
 ip route 40.1.1.0 255.255.255.0 eth1 track 10 state is [up]
 ip route 50.1.1.0 255.255.255.0 20.1.1.4 track 10 state is [up]
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution 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
#
```

Here is the explanation of the "show command" 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.

© 2026 IP Infusion Inc. Proprietary

Table 9. show ip rpf output fields (continued)

Field	Description
RPF type	Different type of RPF like multicast, unicast, MBGP, DVMRP, or static mroutes.
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.

© 2026 IP Infusion Inc. Proprietary

show ipv6 route track-table

Use this command to display IPv4 static routes that are associated to a tracking ID across all VRFs.

Command Syntax

```
show ipv6 route track-table
```

Parameters

None

Default

None

Command Mode

Privileged execution mode and Execution mode

Applicability

This command was introduced in OcNOS version 5.1.

Examples

```
#show ipv6 route track-table
 ipv6 route 5000::/64 eth1 track 10 state is [up]
```

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

Execution mode and Privileged execution 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
#
```

Here is the explanation of the "show command" 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.

Table 10. show ipv6 rpf output fields (continued)

Field	Description
RPF type	Different type of RPF like multicast, unicast, MBGP, DVMRP, or static mroutes.
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.

© 2026 IP Infusion Inc. Proprietary

show rib txlist

Use this command to display internal Routing Information Base (RIB) transaction lists that are pending or currently being transmitted to the Forwarding Information Base (FIB). This command helps monitor routes awaiting installation or synchronization between the routing and forwarding planes. It is mainly used for debugging and operational verification of route processing states.

Command Syntax

```
show rib (txlist | pending-txlist) (ipv4 | ipv6) (vrf WORD|)
```

Parameters

txlist
Displays the list of routes that are currently being transmitted from RIB to FIB.

pending-txlist
Displays the list of routes queued for transmission to the FIB.

ipv4
Displays information for IPv4 routes.

ipv6
Displays information for IPv6 routes.

vrf WORD
Specifies the VRF name for which the RIB transaction information is displayed.

Default

None

Command Mode

Execution mode and Privileged execution mode

Applicability

Introduced before OcNOS version 1.3.

Added the fields (`req_mpls, new_req_mpls, fib_notify_mpls, seq_mpls`) in OcNOS version 7.0.0.

Example

The fields (`req_mpls, new_req_mpls, fib_notify_mpls, seq_mpls`) show MPLS-specific route installation and synchronization details between RIB and FIB.

```
#show rib txlist ipv4
    IPv4:
    0x02739a6c: 127.0.0.0/8        req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
    0x02810cdc: 1.2.0.0/24         req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(5)
    0x0281071c: 1.1.1.1/32         req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(5)
    0x028d9c5c: 2.2.2.2/32         req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(4)
```

© 2026 IP Infusion Inc. Proprietary

```
    0x028da42c: 11.11.11.11/32     req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
    0x028dab2c: 12.12.12.12/32     req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
    0x028daf7c: 13.13.13.13/32     req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
    0x028db0cc: 9.9.9.9/32         req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
    0x028db32c: 10.10.10.10/32     req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
    MARKER MPLS*
    MARKER*

#show rib pending-txlist ipv4 vrf VRF3
IPv4 pending:
0x02634680: 192.0.5.0/24       req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x02634830: 192.0.6.0/24       req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x026349e0: 192.0.7.0/24       req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x02634bb0: 192.0.8.0/24       req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x02652490: 192.0.9.0/24       req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x02652660: 192.0.10.0/24      req(1) new_req(0) fib_notify(0) seq(0) lock(3)
```

© 2026 IP Infusion Inc. Proprietary

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 execution mode and Execution 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 ldp
show running-config router ospf
show running-config router rip
show running-config router rsvp
show running-config router vrrp
```

Parameters

bgp
Display Border Gateway Protocol (BGP) information.

isis
Display Intermediate System to Intermediate System (IS-IS) information.

ldp
Display Label Distribution Protocol (LDP) information.

ospf
Display Open Shortest Path First (OSPF) information.

rip
Display Routing Information Protocol (RIP) information.

rsvp
Display Resource Reservation Protocol (RSVP) information.

vrrp
Display Virtual Router Redundancy Protocol (VRRP) information.

Default

None

Command Mode

Privileged execution mode, Configure mode, Route 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
!
```

© 2026 IP Infusion Inc. Proprietary

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 execution 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

Execution mode and Privileged execution 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
```

© 2026 IP Infusion Inc. Proprietary

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 execution 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
```

© 2026 IP Infusion Inc. Proprietary

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 section describes the commands used to manage the IP SLA for ICMP echo.

clear ip sla statistics 319
frequency 320
icmp-echo 321
ip sla 322
ip sla schedule 323
show ip sla statistics 324
show ip sla summary 326
show running-config ip sla 327
threshold 328
timeout 329

© 2026 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

None

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced in OcNOS version 5.0.

Example

```
#clear ip sla statistics 1
```

© 2026 IP Infusion Inc. Proprietary

frequency

Use this command to configure the frequency or 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.

Example

```
#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
```

© 2026 IP Infusion Inc. Proprietary

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

None

Command Mode

IP SLA mode (config-ip-sla)

Applicability

This command was introduced in OcNOS version 5.0.

Example

```
#configure terminal
(config)#ip sla 1
(config-ip-sla)#icmp-echo ipv4 10.12.28.1 source-interface xe1
(config-ip-sla-echo)#
```

© 2026 IP Infusion Inc. Proprietary

ip sla

Use this command to create an IP SLA instance. One instance maps to a single SLA operation. Users 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

None

Command Mode

Configure 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

WORD
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

None

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 5.0.

Example

```
#configure terminal
(config)#ip sla schedule 1 time-range t1 vrf v1
```

© 2026 IP Infusion Inc. Proprietary

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

None

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced in OcNOS version 5.0.

Example

```
#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
```

Here is the explanation of the "show command" output fields.

Table 11. show ip sla statistics fields

Field	Description
IP SLA ID	IP SLA Identifier (1-65535)
Start Time	Measurement start time

Table 11. show ip sla statistics fields (continued)

Field	Description
Elapsed time(milli sec)	Time taken to complete the measurement in milliseconds
Packets Sent	Number of packet sent
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

None

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced in OcNOS version 5.0.

Example

```
#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
```

Here is the explanation of the "show command" output fields.

Table 12. 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

© 2026 IP Infusion Inc. Proprietary

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

None

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced in OcNOS version 5.0.

Example

```
#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
```

© 2026 IP Infusion Inc. Proprietary

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.

Example

```
#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
```

© 2026 IP Infusion Inc. Proprietary

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.

Example

```
#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
```

© 2026 IP Infusion Inc. Proprietary

Object Tracking Commands

This section describes the Object Tracking commands:

delay up down 331
object-tracking 332
show running-config track 334
show track 335
show track <1-500> 336
show track summary 337
track ip sla reachability 338

© 2026 IP Infusion Inc. Proprietary

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

None

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)#
```

© 2026 IP Infusion Inc. Proprietary

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

Introduced in OcNOS version 6.4.1.

Example

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

```
OcNOS(config)#interface 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
```

© 2026 IP Infusion Inc. Proprietary

```
OcNOS(config-if)#commit
OcNOS(config-if)#exit
```

© 2026 IP Infusion Inc. Proprietary

show running-config track

Use this command to display object tracking running configuration alone.

Command Syntax

```
show running-config track
```

Parameters

None

Default

None

Command Mode

Privileged execution mode and Execution 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#
```

show track

Use this command to display Sham link information.

Command Syntax

```
show track
```

Parameters

None

Default

None

Command Mode

Execution mode

Applicability

This command is introduced in OcNOS version 5.1.

Example

```
OcNOS#show 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

None

Command Mode

Privileged execution mode and Execution mode

Applicability

This command is introduced in OcNOS version 5.1.

Example

```
OcNOS#show 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

None

Default

None

Command Mode

Privileged execution mode and Execution mode

Applicability

This command is introduced in OcNOS version 5.1.

Example

```
OcNOS#show track summary
Object Tracking Summary
ID       Type        Type-Identifier       State
-----------------------------------------------------
1      ip-sla             1                DOWN
2      ip-sla             2                DOWN
OcNOS#
```

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>
Identifier for IP SLA association with tracking object

Command Mode

Configure 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
```

© 2026 IP Infusion Inc. Proprietary

Route-Map Commands

This section describes the Route-Map commands:

continue 341
match as-path 343
match community 344
match extcommunity 345
match interface 346
match ip address 347
match ip address prefix-list 348
match ip next-hop 349
match ip next-hop prefix-list 350
match ip peer 351
match ipv6 address 352
match ipv6 address prefix-list 353
match ipv6 next-hop 354
match ipv6 next-hop prefix-list 355
match ipv6 peer 356
match metric 357
match origin 358
match route-map 359
match route-type 361
match tag 362
route-map 363
set aggregator 364
set as-path 365
set atomic-aggregate 366
set comm-list 367
set community 368
set dampening 370
set extcommunity 371
set interface null0 373
set ip next-hop 374
set ipv6 next-hop 376
set ip next-hop self 377
set large-community 378
set level 380
set local-preference 381
set metric 382

© 2026 IP Infusion Inc. Proprietary

set metric-type 383
set origin 384
set originator-id 385
set tag 386
set vpnv4 next-hop 387
set weight 388
show route-map 389
show running-config route-map 390

© 2026 IP Infusion Inc. Proprietary

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`.

Notes:

- Configuration update from `continue` to `continue <sequence number>` or vice versa is not supported. To switch between these command formats, users must first remove the existing configuration and then apply the new one.
- If the target sequence in the continue statement does not exist, the route-map behaves as if there was no continue statement at all and processing stops after the current sequence.

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

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 to match routes based on BGP standard or expanded community values.

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: Only one `match community` command is permitted per route-map sequence. Configuring a new one replaces the existing entry. This command applies only to BGP routes.

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>
Specifies standard community list number.

<100-199>
Specifies expanded community list number.

WORD
Specifies named community list (standard or expanded).

exact-match
(Optional) Requires the route's communities to exactly match those in the list.

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)#match community mylist
```

© 2026 IP Infusion Inc. Proprietary

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: Only one `match extcommunity` command is permitted per route-map sequence. Configuring a new one replaces the existing entry. This command applies only to BGP routes.

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>
Specifies standard community list number.

<100-199>
Specifies expanded community list number.

WORD
Specifies named community list (standard or expanded).

exact-match
(Optional) Requires the route's communities to exactly match those in the list.

Default

None

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 to disable this function.

Note: The route-map with `match ip address prefix-list WORD` only applies to IPv4 prefixes. When users apply it to IPv6 prefixes, it does not perform a match. Instead, the route-map is treated as RMAP_PERMIT.

Command Syntax

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

Parameter

WORD

IP prefix list name.

Default

None

Command Mode

Route map mode

Applicability

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 differ from what is indicated by the routing protocols. Setting policies allows packets to follow different routes depending on their length or content. Packet forwarding based on configured policies overrides packet forwarding rules specified in routing tables.

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

Notes:

- This command is valid for BGP, OSPFv3, and RIPng only.
- The route-map with `match ipv6 address prefix-list WORD` only applies to IPv6 prefixes. When users apply it to IPv4 prefixes, it does not perform a match. Instead, the route-map is treated as `RMAP_PERMIT`.

Command Syntax

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

Parameter

WORD

IPv6 access list name.

Default

None

Command Mode

Route map mode

Applicability

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 the IPv6 next-hop against entries in a prefix-list. The route defined by policies may differ from the route chosen by routing protocols. Policies allow packets to follow different paths based on their length or content. Packet forwarding determined by configured policies overrides the forwarding decisions made by routing tables. This command applies to BGP and IS-IS.

Note: In 6VPE deployments, the match ipv6 next-hop does not apply to the customer edge (CE) learned IPv6 prefix next-hop because the export map is evaluated only during exporting to the VPNv6 peer. The match condition is only effective when using the provider edges (PEs) IPv4-mapped IPv6 address as the next-hop. For details, see the **`6VPE Configuration`** section in the *OcNOS MPLS Guide*.

Use the `no` parameter with this command to delete the entries in the prefix-list that match the IPv6 next-hop.

Command Syntax

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

Parameters

`prefix-list WORD`
Specifies the name of the IPv6 prefix-list used for matching next-hops.

Default

None

Command Mode

Route map mode

Applicability

Introduced before OcNOS version 1.3.

Examples

This example creates a route-map named `rmap1` with sequence number `3` and attaches a condition to match IPv6 next-hops against the prefix-list `nh_prefixes`. Only routes with next-hops allowed by `nh_prefixes` will match this route-map entry.

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

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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-map

Use this command to nest a route-map as a match criteria. The nested route-map entries are executed in the order of sequence ID. The match criteria of the nested route-map entry are evaluated, then if a match occurs and a permit action is specified in the entry, its set actions are executed and a “match” result is returned to the parent route-map, which then continues its evaluations. If a match occurs and a deny action is specified in the entry, the set actions are skipped, and a no-match result is returned to the parent route-map. If the match criteria are not satisfied in any of the nested route-map entries, a no-match is returned to the parent route-map.

Use the `no` parameter with this command to remove a nested route-map.

Notes: This command applies only to BGP.

A maximum of 8 route-maps can be nested within a route-map instance at the same time and more than 4 levels of nesting are not allowed.

A route-map needs to exist to be nested.

Command Syntax

```
match route-map WORD
no match route-map WORD
```

Parameters

WORD
Name for the route-map (Max size 64).

Default

By default, match route-map is disabled

Command Mode

Route map mode

Applicability

This command was introduced in OcNOS version 7.0.0.

Examples

```
#configure terminal
```

```
(config)#route-map CHILD permit 10
(config-route-map)#exit
(config)#route-map PARENT permit 10
(config-route-map)#match route-map CHILD
```

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
```

© 2026 IP Infusion Inc. Proprietary

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)#
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

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 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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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>
set extcommunity color <1-4294967295>
no set extcommunity rt (.AA:NN|) (additive|)
no set extcommunity soo (.AA:NN|)
no set extcommunity cost (igp|pre-bestpath|) <0-255> <0-4294967295>
no set extcommunity color <1-4294967295>
```

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.

`color`
Color value extended community

Default

None

© 2026 IP Infusion Inc. Proprietary

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

#configure terminal
(config-route-map)#route-map rmap1 permit 3
(config-route-map)#set extcommunity color 101
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

set ip next-hop

Use this command to set the specified next-hop value and configure IPv4 nexthop as tracked object to verify its availability.

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

Command Syntax

```
set ip next-hop A.B.C.D | verify-availability track <1-500> | peer-address |self
no set ip next-hop (A.B.C.D|)
no set ip next-hop peer-address
no set ip next-hop self
```

Parameter

A.B.C.D
IP address of the next-hop.

verify-availability track <1-500>
Defines a tracking object (numbered 1 to 500) to verify-availability of IPv4 nexthop.

Note: Depends on “IP-SLA” configurations.

peer-address
Use peer address (for Outbound BGP only)

self
Self (for BGP LU only)

Note: The **set ip next-hop** options are mutually exclusive. Only one option can be configured in a route-map at any given time. Configuring more than one option is not permitted.

Default

Disabled

Command Mode

Route map mode

Applicability

This command was introduced before OcNOS version 1.3 and modified in OcNOS version 6.6.1 to include the parameter `verify-availability track`.

Examples

The following example is for configuring a route-map to set the next-hop IP address, without availability trackings:

```
#configure terminal
```

```
(config)#route-map mymap permit 3
(config-route-map)#set ip next-hop 10.10.0.67
```

The following example is for configuring a route-map to set the next-hop IPv4 address, with availability trackings using a tracking object:

```
#configure terminal
(config)#route-map mymap permit 3
(config-route-map)#set ip next-hop 10.10.0.67 verify-availability track 7
```

© 2026 IP Infusion Inc. Proprietary

set ipv6 next-hop

Use this command to set a next hop-address and configure IPv6 nexthop as tracked object to verify its availability.

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 | verify-availability track <1-500>
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.

`verify-availability track <1-500>`
Defines a tracking object (numbered 1 to 500) to verify-availability of IPv6 nexthop.

Note: Depends on “IP-SLA” configurations.

`local`
Local IPv6 address of the nexthop.

Default

Disabled

Command Mode

Route map mode

Applicability

This command was introduced before OcNOS version 1.3 and modified in OcNOS version 6.6.1 to include the parameter `verify-availability track`.

Examples

The following example is for configuring a route-map to set the next-hop IP address, without availability trackings:

```
#configure terminal
(config)#route-map rmap1 permit 3
(config-route-map)#set ipv6 next-hop local fe80::203:47ff:fe97:66dc
```

The following example is for configuring a route-map to set the next-hop IPv6 address, with availability trackings using a tracking object:

© 2026 IP Infusion Inc. Proprietary

```
#configure terminal
(config)#route-map mymap permit 3
(config-route-map)#set ipv6 next-hop 1311::2 verify-availability track 8
```

set ip next-hop self

Use this command to set next hop self for IPV4 BGP-LU neighbors.

Use “no” form of this command to not set next hop self.

Command Syntax

```
set ip next-hop self
no set ip next-hop self
```

Parameters

None

Command Mode

Route map mode mode

Applicability

This command is introduced in OcNOS version 7.0.0

Examples

```
ip prefix-list BCOM-IP
   seq 5 permit 25.4.4.0/24 eq 24

route-map BCOM-RM permit 2
   match ip address prefix-list BCOM-IP
   set ip next-hop self



router bgp 26
..
  address-family ipv4 labeled-unicast
  neighbor 27.27.27.27 route-map BCOM-RM out
  ..
```

© 2026 IP Infusion Inc. Proprietary

set large-community

Use this command to set the large community values in routing policies.

Command Syntax

```
set large-community (.AAaa:NNnn:ZZzz) (additive|)
no set large-community (.AAaa:NNnn:ZZzz) (additive|)
```

Parameters

.AAaa:NNnn:ZZzz
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 is enabled by default for peers.

Command Mode

Route map mode

Applicability

This command was introduced in OcNOS version 6.1.0. Introduced a new parameter, `additive`, and updated the maximum character length for large community values to `255` in OcNOS version 6.6.1.

Examples

Setting Multiple Large Communities

```
OcNOS(config)#route-map R1 permit 10
OcNOS(config-route-map)#set large-community 2:2:2 9:9:9 3:3:3
OcNOS(config-route-map)#commit
```

The below configuration adds two large community attributes `132:1:132` and `132:1:132`, to the matching routes. The `additive` option ensures these communities are added without replacing any existing ones.

```
OcNOS(config-route-map)#set large-community 132:1:132 132:1:132 additive
OcNOS(config-route-map)#commit
```

Setting a Single Large Community

```
OcNOS(config-route-map)#set large-community 1:2:3
OcNOS(config-route-map)#commit
```

- Verify the configured lists using the `show` command.

```
OcNOS#show running-config route-map
!
route-map rmap permit 10
 match large-community EXPTEST
 set large-comm-list 100 delete
!
```

```
route-map rmap1 permit 10
 match large-community 1
 set large-community 1:2:3
!
route-map rmap2 permit 10
 match large-community 100
 set large-comm-list 104 delete
!
route-map rmap3 permit 10
 match large-community STDTEST exact-match
!

OcNOS#show route
route-map  router-id
OcNOS#show route-map
route-map rmap, permit, sequence 10
  Match clauses:
    large-community: EXPTEST
  Set clauses:
    large-comm-list 100 delete
route-map rmap1, permit, sequence 10
  Match clauses:
    large-community: 1
  Set clauses:
    large-community 1:2:3
route-map rmap2, permit, sequence 10
  Match clauses:
    large-community: 100
  Set clauses:
    large-comm-list 104 delete
route-map rmap3, permit, sequence 10
  Match clauses:
    large-community: STDTEST exact-match
  Set clauses:
```

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

set vpnv4 next-hop

Use this command to set a VPNv4 next-hop address.

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.

Note: This command is valid for BGP only.

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

Command Syntax

```
set vpnv4 next-hop A.B.C.D
no set vpnv4 next-hop (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 vpnv4 next-hop 6.6.6.6
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution 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
#
```

Here is the explanation of the "show command" output fields.

Table 13. 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.

© 2026 IP Infusion Inc. Proprietary

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 execution 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
!
```

© 2026 IP Infusion Inc. Proprietary

Source Interface Commands

This section 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 the table below:

Table 14. Source interface protocols and port numbers

Protocol	Default port number
Tacacs+	49
Radius	1812 and 1813
Snmp	161 and 162
Ntp	123
Syslog	514

This section contains these commands:

ip source-interface 392
ipv6 source-interface 394
show ip source-interface detail 396
show ipv6 source-interface detail 398
show running-config ip source-interface 400
show running-config ipv6 source-interface 401

© 2026 IP Infusion Inc. Proprietary

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.

Notes:

- 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 (page 392) or ipv6 source-interface (page 394) is used.
- NTP is supported only on default port.
- OcNOS supports configuring multiple ports for the same protocol in the `ip source-interface` command.

Command Syntax

```
ip source-interface IFNAME (tacacs+|ntp|snmp|syslog|radius) (port (123|162|1812|49|514|<1025-
65535>)|)(vrf (NAME|management)|)
no ip source-interface IFNAME (tacacs+|ntp|snmp|syslog|radius) (port (49|123|162|514|1812|<1025-
65535>)|) (vrf (NAME|management)|)
```

Parameters

IFNAME
Interface name (lo or physical interface)

tacacs+
Terminal Access Controller Access Control System

tacacs+ 49
Default source interface protocol port number for TACACS+

ntp
Network Time Protocol

ntp 123
Default source interface protocol port number for NTP

snmp
Simple Network Management Protocol

snmp 162
Default source interface protocol port number for SNMP

syslog
syslog

syslog 514
Default source interface protocol port number for Rsyslog

radius
Remote Authentication Dial-In User Service

© 2026 IP Infusion Inc. Proprietary

radius 1812
Default source interface protocol port number for RADIUS

<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

None

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 xe12 ntp vrf vrf1
(config)#ip source-interface xe2 radius port 1025
(config)#ip source-interface xe3 syslog port 65535 vrf management
```

The following example shows how to configure multiple source interface protocol port number:

```
OcNOS#configure terminal
OcNOS(config)#ip source-interface lo syslog port 6666
OcNOS(config)#ip source-interface lo syslog port 7777
OcNOS(config)#ip source-interface lo syslog port 8888
OcNOS(config)#ip source-interface lo syslog port 9999
OcNOS(config)#commit
```

© 2026 IP Infusion Inc. Proprietary

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 (123|162|1812|49|514|<1025-
65535>)|) (vrf (NAME|management)|)
no ipv6 source-interface IFNAME (tacacs+|ntp|snmp|syslog|radius) (port (49|123|162|514|1812|<1025-
65535>)|) (vrf (NAME|management)|)
```

Parameters

IFNAME
Interface name (lo or physical interface)

tacacs+
Terminal Access Controller Access Control System

tacacs+ 49
Default source interface protocol port number for TACACS+

ntp
Network Time Protocol

ntp 123
Default source interface protocol port number for NTP

snmp
Simple Network Management Protocol

snmp 162
Default source interface protocol port number for SNMP

syslog
syslog

syslog 514
Default source interface protocol port number for Rsyslog

radius
Remote Authentication Dial-In User Service

radius 1812
Default source interface protocol port number for RADIUS

<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

None

© 2026 IP Infusion Inc. Proprietary

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

```
OcNOS#configure terminal
OcNOS(config)#ipv6 source-interface xe12 ntp vrf vrf1
OcNOS(config)#ipv6 source-interface xe2 radius port 1025
OcNOS(config)#ipv6 source-interface xe3 syslog port 65535 vrf management
```

The following example shows how to configure multiple source interface protocol port number:

```
OcNOS#configure terminal
OcNOS(config)#ipv6 source-interface lo tacacs+ port 1111
OcNOS(config)#ipv6 source-interface lo tacacs+ port 2222
OcNOS(config)#ipv6 source-interface lo tacacs+ port 3333
OcNOS(config)#ipv6 source-interface lo tacacs+ port 4444
OcNOS(config)#commit
OcNOS(config)#
```

© 2026 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

Default

None

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

The following example shows the output using different protocols with multiple port numbers:

```
#show ip source-interface detail
Source-Interface Detailed Information
=====================================
  Protocol : tacacs+
  Interface : lo
  Port : 6666
  Address : 1.1.1.1
  Status : Active
  VRF Name : Default

  Protocol : radius
  Interface : lo
  Port : 7777
  Address : 1.1.1.1
  Status : Active
  VRF Name : Default

  Protocol : ntp
  Interface : xe12
  Port : 8888
  Address : 1.1.1.1
  Status : Active
  VRF Name : Management

  Protocol : syslog
  Interface : lo
  Port : 9999
  Address : 0.0.0.0
  Status : Active
  VRF Name : default
```

© 2026 IP Infusion Inc. Proprietary

```
 Protocol : ntp
  Interface : xe12
  Port : 123
  Address : 1.1.1.1
  Status : Active
  VRF Name : vrf1
```

The following example shows the output using the same protocol with multiple port numbers:

```
OcNOS#show ip source-interface detail
Source-Interface Detailed Information
=====================================
  Protocol : syslog
  Interface : lo
  Port : 6666
  Address : 0.0.0.0
  Status : Active
  VRF Name : default

  Protocol : syslog
  Interface : lo
  Port : 7777
  Address : 0.0.0.0
  Status : Active
  VRF Name : default

  Protocol : syslog
  Interface : lo
  Port : 8888
  Address : 0.0.0.0
  Status : Active
  VRF Name : default

  Protocol : syslog
  Interface : lo
  Port : 9999
  Address : 0.0.0.0
  Status : Active
  VRF Name : default
```

Here is the explanation of the "show command" output fields.

Table 15. Output fields

Field	Description
Protocol	tacacs+, ntp, snmp, syslog, or radius
Interface	Interface name (lo or physical interface)
Port	Source interface protocol port number
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

Default

None

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

The following example shows the output using different protocols with multiple port numbers:

```
#show ipv6 source-interface detail
Source-Interface Detailed Information
=====================================
  Protocol : tacacs+
  Interface : lo
  Port : 6666
  Address : ::1
  Status : Active
  VRF Name : Default

  Protocol : radius
  Interface : lo
  Port : 7777
  Address : ::1
  Status : Active
  VRF Name : Default

  Protocol : ntp
  Interface : xe12
  Port : 7777
  Address : 2001::1
  Status : Active
  VRF Name : vrf1
```

The following example shows the output using the same protocol with multiple port numbers:

```
OcNOS#show ipv6 source-interface detail
Source-Interface Detailed Information
=====================================
```

© 2026 IP Infusion Inc. Proprietary

```
Protocol : tacacs+
Interface : lo
Port : 1111
Address : ::
Status : Active
VRF Name : default

Protocol : tacacs+
Interface : lo
Port : 2222
Address : ::
Status : Active
VRF Name : default

Protocol : tacacs+
Interface : lo
Port : 3333
Address : ::
Status : Active
VRF Name : default

Protocol : tacacs+
Interface : lo
Port : 4444
Address : ::
Status : Active
VRF Name : default
```

Here is the explanation of the "show command" output fields.

Table 16. Output fields

Field	Description
Protocol	tacacs+, ntp, snmp, syslog, or radius
Interface	Interface name (lo or physical interface)
Port	Source interface protocol port number
Address	IP address
Status	Whether active or inactive
VRF Name	Virtual Routing and Forwarding name

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

Default

None

Command Mode

Execution mode and Privileged execution 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
ip source-interface lo syslog port 6666
ip source-interface lo syslog port 7777
ip source-interface lo syslog port 8888
ip source-interface lo syslog port 9999
```

© 2026 IP Infusion Inc. Proprietary

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

Default

None

Command Mode

Execution mode and Privileged execution 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
ip source-interface lo syslog port 6666
ip source-interface lo syslog port 7777
ip source-interface lo syslog port 8888
ip source-interface lo syslog port 9999
```

© 2026 IP Infusion Inc. Proprietary

BORDER GATEWAY PROTOCOL CONFIGURATION

BGP 408
Enable BGP Routers in the Same Autonomous System 408
Topology 408
Validation 409
Enable BGP Between Different Autonomous Systems 409
Route-Map 411
Route-map Continue 413
Route-map Nesting 420
Route Reflector 424
Multiple Route Reflectors 428
BGP Confederations 438
Dynamic BGP Peering 450
Enable eBGP Multihop 479
Enable Peer Groups 482
Route Redistribution in BGP 493
Add Multiple Instances of the Same Autonomous System 494
Remove the Multi-Exit Disc Attribute from Update Messages 497
Removing Sent and Received MED values 498
BGP Four-Byte Autonomous System 500
4-Octet ASN Capability Enabled on R1 and R2 500
4-Octet ASN Capability Enabled on R1 and Disabled on R2 501
BGP Extended Community Attribute 504
Nexthop Tracking 521
Nexthop Tracking Delay Timer 523
BGP Distance 530
BGP Weight per Peer 533
OSPF as PE-CE Protocol for VPNs 536
BGP Multipath for IPv4 539
Multipath eBGP 543
Multipath eiBGP 546
BGP AS-PATH Multipath-relax 550
BGP FIB Install (Selective Route Download) 553
Route Target Constraint 559
BGP Best Path Selection Process 569
BGP Dampening 570

© 2026 IP Infusion Inc. Proprietary

BGP Authentication 576
BGP Unnumbered 579
BGP Blackhole Community Attribute 587
Accumulated Interior Gateway Protocol 590
BGP Graceful Restart Configuration 601
BGP Automated Dynamic Route Policy Update 608
Overview 608
Prerequisites 609
Configuring BGP auto-policy-soft-reset 609
Implementation Examples 617
CLI Commands 618
Glossary 620
BGP4+ 621
Enable iBGP Peering Using a Global Address 621
Configuration 621
Validation 622
Enable iBGP Peering Using Link-local Address 624
Configuration 624
Validation 625
Enable eBGP Peering Between Different Autonomous Systems 626
Configuration 627
Validation 628
Route-Map 629
Configuration 630
Validation 631
Route Reflector 633
Configuration 633
Validation 635
Confederations 637
Configuration 637
Validation 640
BGP4+ Graceful Restart 642
Configuration 642
Validation 643
Configure BGP4+ Distance 647
Configuration 648
Validation 649
BGP4+ Graceful Reset 652
Configuration 652
Validation 653

VPNv4 NLRI with IPv6 Nexthop 656
Configuration 656
Validation 659
BGP Labeled Unicast 662
BGP Labeled Unicast as Transport Configuration 663
BGP Labeled Unicast with Seamless MPLS Configuration 675
BGP Labeled Unicast with Inter-AS Configuration 695
BGP Peer Groups for Address-Family IPv4 Labeled Unicast 709
BGP Labeled Unicast with Implicit Null Label for Local Routes 720
Overview 720
Prerequisites 720
Topology 721
Configuration 721
Validation 723
Glossary 728
BGP MPLS Next Hop Tracking Configuration 729
Configuration 729
Validation 737
BGP Labeled Unicast Next Hop Self in Route-Map 742
Overview 742
Feature Characteristics 742
Limitations 742
Prerequisites 742
Configuration 742
Commands 802
BGP Additional Paths for VPNv4 803
Overview 803
Topology 804
Initial Configuration 804
RR 831
PE1 832
PE2 832
PE3 832
PE4 833
RR 833
IPV6 LABELED-UNICAST Prefix Advertised to All 6pe Peers With All Addpaths: 836
PE1 836
PE2 837
PE3 837
PE4 838

CE1 839
RR 839
Validation 839
RR 839
PE1 840
PE2 841
PE3 841
PE4 842
RR 843
Validation 843
RR 843
PE1 844
PE2 845
PE3 845
PE4 846
RR 846
Validation 847
PE1 847
RR 847
Validation 848
PE1 848
RR 848
Validation 849
PE1 849
RR 850
Validation 850
PE1 850
CE1 850
CE2 851
PE1 851
PE2 852
RR 852
PE3 854
CE3 854
Additional Paths for 6VPE 854
BGP Labeled Unicast IPv4 with Add Path 887
BGP4+ Additional Paths Configuration 971
Overview 971
Topology 973
Additional Paths at the Global Level 978

© 2026 IP Infusion Inc. Proprietary

Additional Paths at the Neighbor Level ... 979
Additional Paths Send and Receive at Address-family level ... 984
Additional Paths at the Neighbor Level ... 985
Selection of all Additional Paths at the Address-family Level ... 991
Selection of all Additional Paths at the Neighbor Level ... 992
Selection of Best 2 Additional Paths at AF Level ... 994
Selection of Best 2 Additional Paths at the Neighbor Level ... 995
Selection of Best 3 Additional Paths at the AF Level ... 997
Selection of Best 3 Additional Paths at the Neighbor Level ... 998
BGP Flowspec Support for IPv4 ... 1000
BGP Flowspec IPv4 Configuration ... 1000
BGP Flowspec Configuration With RR ... 1060
TCP MSS configuration for BGP neighbors ... 1128
Overview ... 1128
Prerequisites ... 1129
Configuration ... 1129
Validation ... 1131
New CLI Commands ... 1132
Abbreviations ... 1133
Glossary ... 1134
BGP ORF Prefix-List VPNV4 Address ... 1135
Overview ... 1135
Configuration ... 1135
Topology ... 1135
Glossary ... 1146
BGP Bogon Prefix Filtering IPv4 ... 1148
Overview ... 1148
BGP Bogon Prefix Filtering Configuration ... 1148
BGP Bogon Prefix Filtering Command ... 1151
Glossary ... 1152
BGP Large Communities ... 1153
Overview ... 1153
BGP Large Communities Configuration ... 1153
Implementation Example ... 1171
BGP Large Communities Commands ... 1171
Troubleshooting ... 1180
Glossary ... 1180
Hide the Remote AS using the neighbor local-as Command ... 1181
Overview ... 1181
Configuration ... 1181

© 2026 IP Infusion Inc. Proprietary

Validation 1182
Revised CLI Commands 1184
BGP Aggregate-Address in VRF Address-family 1186
Overview 1186
Prerequisites 1186
Configuration for Aggregate-Address in VRF Address-Family with IPv4 1186
Running Configuration with IPv4 1190
Configuration for Aggregate-Address in VRF Address-Family with IPv6 1193
Running Configuration with IPv6 1195
BGP RPKI-Based Route Validation 1198
Overview 1198
Feature Characteristics 1198
Benefits 1198
Configuration 1199
Implementation Examples 1203
Commands 1203
Glossary 1217
Recursive Next-hop Resolution with MPLS Next-hop 1219
Overview 1219
Limitation 1219
Configuration 1219
Implementation Example 1232
Revised Commands 1232
Troubleshooting 1233
Glossary 1235

© 2026 IP Infusion Inc. Proprietary

BGP

This section contains basic Border Gateway Protocol configuration examples.

Enable BGP Routers in the Same Autonomous System

Figure 21 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 21. Routers in the Same Autonomous System

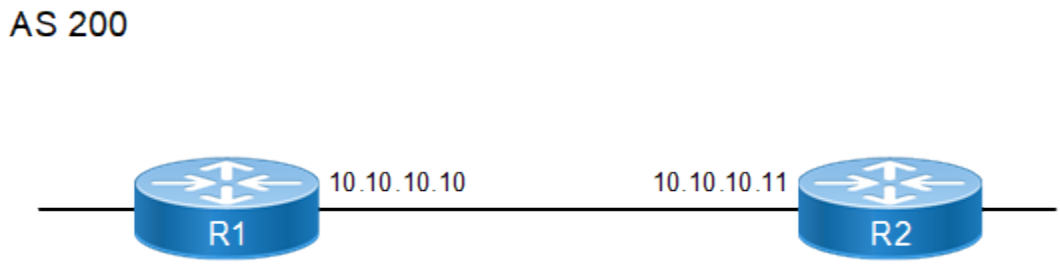

R1

`#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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

R2

`#configure terminal`	Enter configure mode.
`(config)#router bgp 200`	Define the routing process. The number `200` specifies the AS number of `R2`.

© 2026 IP Infusion Inc. Proprietary

`(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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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 22. Routers in Different Autonomous Systems

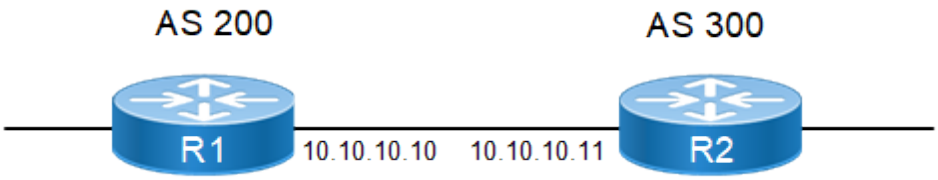

R1

`#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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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
```

© 2026 IP Infusion Inc. Proprietary

```
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
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.

Note: Route maps enable filtering and modification of BGP routes using 'match' and 'set' clauses. They can be applied to a BGP peer or peer-group in either the inbound or outbound direction to influence received or advertised routes. However, for route-map updates to take effect, a manual command such as clear ip bgp A.B.C.D`<neighbor> soft` must be executed for each BGP peer or peer-group address-family. This manual step is required in all software releases except for release 6.6.x and later, provided that the 'bgp auto-policy-soft-reset enable' feature is configured.

© 2026 IP Infusion Inc. Proprietary

Topology

Figure 23. Configure Route-Map

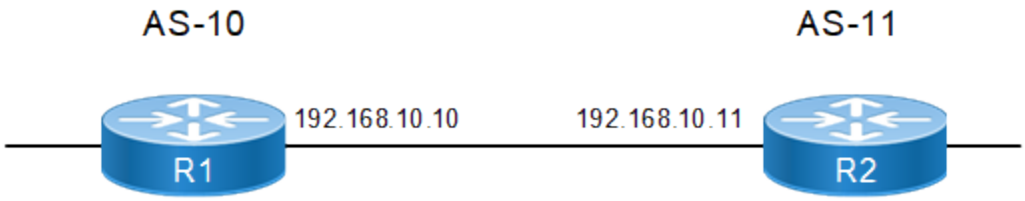

R1

`#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)#commit`	Commit the candidate configuration to the running configuration.
`(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

© 2026 IP Infusion Inc. Proprietary

	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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

R2

(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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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-map Continue

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 24. Route-map continue

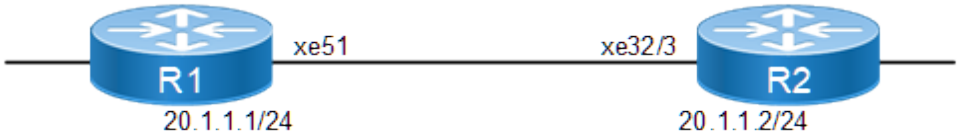

Configure Route-map continue on R1

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)#commit`	Commit the candidate configuration to the running configuration
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#ip prefix-list PF1`	Configure IP prefix-list PF1
`R1(config-ip-prefix-list)#seq 5 permit 10.1.0.0/16`	Create an access rule to permit IP packets with

© 2026 IP Infusion Inc. Proprietary

`le 32`	Maximum prefix length match
R1(config-ip-prefix-list)#exit	Exit interface mode
R1(config-ip-prefix-list)#ip prefix-list P1	Configure IP prefix list P1
`R1(config-ip-prefix-list)#seq 10 permit 10.1.1.0/24 le 32`	Create an access rule to permit IP packets with Maximum prefix length match
R1(config-ip-prefix-list)#exit	Exit interface mode
R1(config-ip-prefix-list)#ip prefix-list P2	Configure IP prefix list P2
`R1(config-ip-prefix-list)#seq 15 permit 10.1.2.0/24 le 32`	Create an access rule to permit IP packets with Maximum prefix length match
R1(config-ip-prefix-list)#exit	Exit interface mode
R1(config-ip-prefix-list)#ip prefix-list P3	Configure IP prefix list P3
`R1(config-ip-prefix-list)#seq 20 permit 10.1.3.0/24 le 32`	Create an access rule to permit IP packets with Maximum prefix length match
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration
`R1(config-if)#exit`	Exit interface 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 65535`	Set weight as 65535
R1(config-route-map)#exit	Exit Route map mode
`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)#exit	Exit Route map mode
`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)#exit	Exit Route map mode

© 2026 IP Infusion Inc. Proprietary

`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-if)#commit`	Commit the candidate configuration to the running configuration
`R1(config-route-map)#exit`	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
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 unicast	Enter address-family mode
R1(config-router-af)#neighbor 20.1.1.2 activate	Activate neighbor
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)#exit-address-family	Exit address-family mode
R1(config-router-af)#commit	Exit router bgp mode
R1(config-router-af)#exit	Commit the candidate configuration to the running configuration

R2

`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)#commit`	Commit the candidate configuration to the running configuration
`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)#commit`	Commit the candidate configuration to the running configuration
`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)#commit`	Commit the candidate configuration to the running configuration
`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

© 2026 IP Infusion Inc. Proprietary

R2(config-if)#commit	Commit the candidate configuration to the running configuration
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)#commit	Commit the candidate configuration to the running configuration
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)#commit	Commit the candidate configuration to the running configuration
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 unicast	Enter address-family mode
R2(config-router-af)#neighbor 20.1.1.1 activate	Activate neighbor
R2(config-router)#redistribute connected	Redistribute the connected routes which are 10 networks here.
R2(config-if)#commit	Commit the candidate configuration to the running configuration
R2(config-if)#exit	Exit interface mode.

Validation

R1

The following provides the R1 validation:

```
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#
```

© 2026 IP Infusion Inc. Proprietary

R2

The following provides the R2 validation:

```
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
```

Route-map Nesting

Route-maps can be used to filter incoming updates from a BGP peer. The `match route-map` command within a route-map allows another route-map to be called as a subroutine. The nested route-map is evaluated, and if it returns permit, it is treated as a match; if it returns deny, it is treated as a no-match. A route-map must exist to be nested.

Topology

In the below example, router R2 sets the metric to 300 for all the routes except those with the prefix 20.1.2.0/23. Router R1 then filters incoming route advertisements by matching metric 300 and matching the nested route-map

CHILD.

In the CHILD route-map, the sequence 10 denies prefix 10.1.1.0/24, while sequence 20 denies prefix 1.1.1.0/24 and permits all other routes. Since 20.1.2.0/23 is the only route from R2 with a metric different from 300, it is denied.

Figure 25. Route-map Nesting

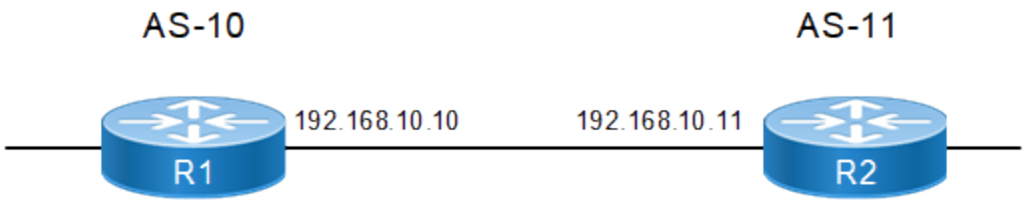

Note: The nested route-map is referred to only by the name, so the first route-map in the sequence with the same name will be evaluated, and if no match is found, the next in the sequence will be evaluated. This continues until a match is found or the sequence ends. A maximum of 8 route-maps can be nested within a route-map instance at the same time. Note that more than 4 levels of nesting are not allowed.

Configure Route-map

R1

`R1#configure terminal`	Enter configure mode.
`R1(config)#ip prefix-list ABC`	Create an entry in the prefix-list. The `ABC` parameter is the name of the prefix-list.
`R1(config-ip-prefix-list)#seq 5 deny 1.1.1.0/24 ge 26 le 28`	The sequence number of the prefix entry is specified as `5`. `deny` specifies that the packets must be rejected. The minimum and maximum prefix lengths to be matched are `26` and `28`.
`R1(config-ip-prefix-list)#seq 10 permit any`	The sequence number of the prefix-list entry is specified as 10. The permit parameter `any` denotes accepting all packets of any length.
`R1(config-ip-prefix-list)#exit`	Exit prefix-list mode.
`R1(config)#ip prefix-list ABC2`	Create an entry in the prefix-list. The `ABC2` parameter is thename of the prefix-list.
`R1(config-ip-prefix-list)#seq 5 permit 10.1.1.0/24`	The sequence number of the prefix entry is specified as `5`. The parameter `permit` specifies that the packets with the prefix 10.1.1.0/24 will be accepted.
`R1(config-ip-prefix-list)#exit`	Exit prefix-list mode
`R1(config)#route-map PARENT permit 10`	Create the PARENT route-map.
`R1(config-route-map)#match route-map CHILD`	Match route-map CHILD as a subroutine. If a match is returned, this match condition is satisfied.

© 2026 IP Infusion Inc. Proprietary

`R1(config-route-map)#match metric 300`	Match only the routes with the metric 300.
`R1(config-route-map)#exit`	Exit Route-map mode and return to Configure mode.
`R1(config)#route-map CHILD deny 10`	Create an entry with the sequence number `10` and a `deny` action for the route-map that is being nested.
`R1(config-route-map)#match ip address prefix-list ABC2`	Set the match criteria based on the prefix-list `ABC2`. As the prefix 10.1.1.0/24 will be a match, and the action of the route-map is “deny”, the routes with the prefix 10.1.1.0/24 will be denied. If an entry does not have a match, the next route-map in the sequence will be evaluated.
`R1(config-route-map)#exit`	Exit Route-map mode and return to Configure mode.
`R1(config)#route-map CHILD permit 20`	Create an entry with the sequence number `20` and a `permit` action for the route-map that is being nested.
`R1(config-route-map)#match ip address prefix-list ABC`	Set the match criteria based on the prefix-list `ABC`. Every prefix that is matched by the prefix-list will be accepted.
`R1(config-route-map)#exit`	Exit Route-map mode and return to Configure mode.
`R1(config)#router bgp 10`	Define the routing process, and establish a TCP session. The number 10 specifies the AS number of R1.
`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.
`R1(config-router)#address-family ipv4 unicast`	Enter address-family ipv4 unicast mode.
`R1(config-router-af)#neighbor 192.168.10.11 activate`	Activate the neighbor in the IPv4 address family.
`R1(config-router-af)#neighbor 192.168.10.11 route-map PARENT in`	Apply a route map to routes. 192.168.10.11 specifies the IP address of the BGP neighbor. The name of the route map is PARENT, and the `in` parameter specifies that the access list applies to incoming advertisements.
`R1(config-router-af)#exit-address-family`	Exit address-family mode.
`R1(config-router)#commit`	Commit the candidate configuration to the running configuration.

R2

`R2(config)#ip prefix-list PLIST`	Create an entry in the prefix-list. The name of the prefix-list is `PLIST`.
`R2(config-ip-prefix-list)#seq 5 deny 20.1.2.0/23`	The sequence number or position of the prefix entry is specified as `5`. `deny` specifies that the packets with the prefix

	20.1.2.0/23 are to be rejected.
`R2(config-ip-prefix-list)#seq 10 permit any`	The sequence number or position of this prefix entry is specified as `10`. The permit parameter `any` specifies that all packets of any length are accepted.
`R2(config-ip-prefix-list)#exit`	Exit the prefix-list mode .
`R2(config)#route-map TEST permit 10`	Enter Route-map mode to set the match operation.
`R2(config-route-map)#match ip address prefix-list PLIST`	Set the match criteria based on the prefix-list PLIST
`R2(config-route-map)#set metric 300`	Configure the set action. If the match condition is satisfied, the metric is set to 300.
`R2(config-route-map)#exit`	Exit Route-map mode and return to Configure mode.
`R2 (config)#interface lo`	Enter loopback interface mode.
`R2(config-if)#ip address 1.1.1.1/27 secondary`	Specify a secondary interface address.
`R2(config-if)#ip address 1.1.2.1/24 secondary`	Specify a secondary interface address
`R2(config-if)#ip address 10.1.1.1/24 secondary`	Specify a secondary interface address
`R2(config-if)#ip address 20.1.1.1/24 secondary`	Specify a secondary interface address
`R2(config-if)#ip address 20.1.2.1/23 secondary`	Specify a secondary interface address
`R2(config-if)#ip address 30.1.1.1/28 secondary`	Specify a secondary interface address
`R2(config-if)#exit`	Exit loopback interface mode.
`R2(config)#router bgp 11`	Define the routing process, and establish a TCP session. The number `11` specifies the AS number of R2.
`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.
`R2(config-router)#address-family ipv4 unicast`	Enter address-family IPv4 unicast mode
`R2(config-router-af)#neighbor 192.168.10.10 activate`	Activate the neighbor in the IPv4 address family.
`R2(config-router-af)#neighbor 192.168.10.10 route-map TEST out`	Apply a route map to routes. `192.168.10.10` specifies the IP address of the BGP neighbor. The `TEST` parameter is the name of the route map, and `out` specifies that the route-map applies to outgoing advertisements.
R2(config-router-af)#network 1.1.1.0/27	Specify the network to be advertised by the BGP routing process.
R2(config-router-af)#network 1.1.2.0/24	Specify the network to be advertised by the BGP routing process.
`R2(config-router-af)#network 10.1.1.0/24`	Specify the network to be advertised by the BGP routing process.
`R2(config-router-af)#network 20.1.1.0/24`	Specify the network to be advertised by the BGP routing process.

© 2026 IP Infusion Inc. Proprietary

`R2(config-router-af)#network 20.1.2.0/23`	Specify the network to be advertised by the BGP routing process.
`R2(config-router-af)#network 30.1.1.0/28`	Specify the network to be advertised by the BGP routing process.
`R2(config-router-af)#exit-address-family`	Exit router mode.
`R2(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

Here is the validation:

```
#sh ip bgp
BGP table version is 2, local router ID is 172.16.0.14
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid,
 > best, i - internal,
              l - labeled, S Stale, x-EVPN
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

    Network          Next Hop            Metric    LocPrf   Weight Path   Ext-Co
lor
*>   1.1.2.0/24       192.168.10.11         300      100       0   11 i        -
*>   20.1.1.0/24      192.168.10.11         300      100       0   11 i        -
*>   30.1.1.0/28      192.168.10.11         300      100       0   11 i        -

Total number of prefixes 3
```

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).

Notes:

- When utilizing BGP redundant route reflectors with BFD, setting the multi-hop timer to 999ms and the multiplier to 3 results in a BGP convergence time of approximately 6 seconds in the event of an aggregator node reload.

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.

© 2026 IP Infusion Inc. Proprietary

Figure 26. BGP Route Reflector

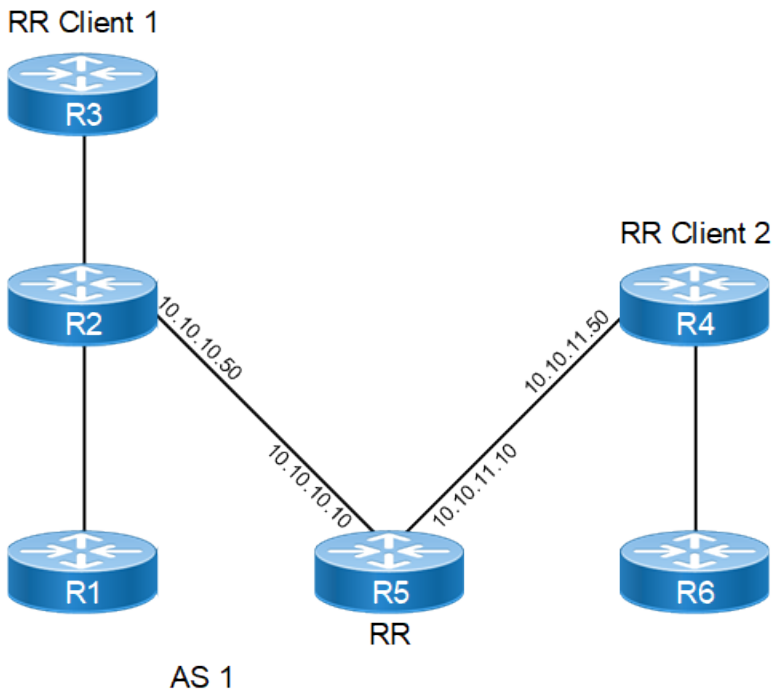

RR (R5)

`#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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

© 2026 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.
`(config)#commit`	Commit the candidate configuration to the running configuration.

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.
`(config)#commit`	Commit the candidate configuration to the running configuration.

Validation

R5

The following provides the R5 validation:

```
#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

The following provides the R3 validation:

```
#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

The following provides the R2 validation:

```
#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.

© 2026 IP Infusion Inc. Proprietary

Topology

Figure 27. eBGP and iBGP Route Reflector Topology

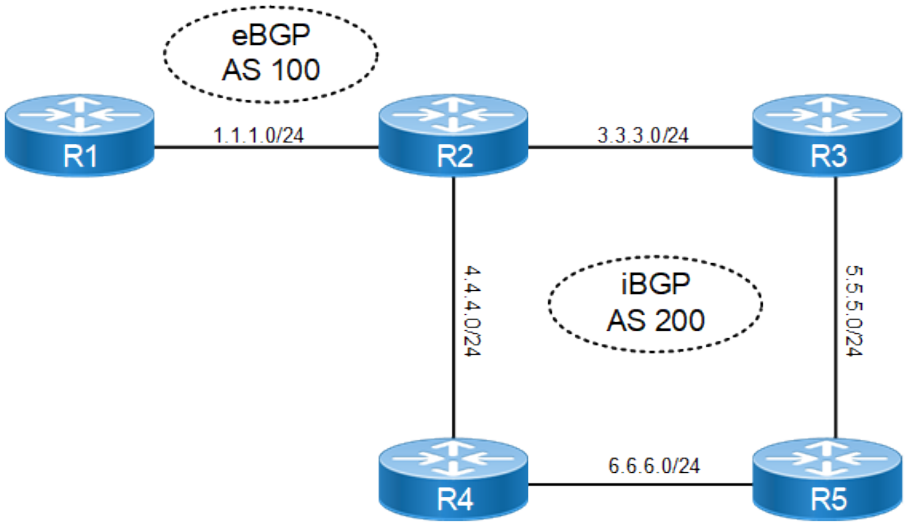

R1

`#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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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

© 2026 IP Infusion Inc. Proprietary

`(config)#interface eth2`	Enter interface mode
`(config-if)#ip address 3.3.3.2/24`	Specify IP address for the interface.
`(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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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

© 2026 IP Infusion Inc. Proprietary

`(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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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.
`(config)#commit`	Commit the candidate configuration to the running configuration.

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).

© 2026 IP Infusion Inc. Proprietary

`(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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

R2

The following provides the R2 validation:

```
#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

The following provides the R1 validation:

```
#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

The following provides the R3 validation:

```
#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

The following provides the R4 validation:

```
#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: ::
```

© 2026 IP Infusion Inc. Proprietary

```
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

The following provides the R5 validation:

```
#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 28. BGP Confederation

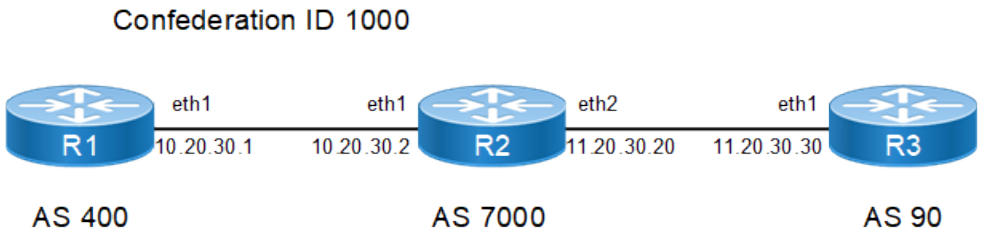

R1

`#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.
`(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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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).

© 2026 IP Infusion Inc. Proprietary

(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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

R2

The following provides the R2 validation:

```
#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

The following provides the R1 validation:

```
#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

The following provides the R3 validation:

```
#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 29. BGP Confederation with Multiple AS

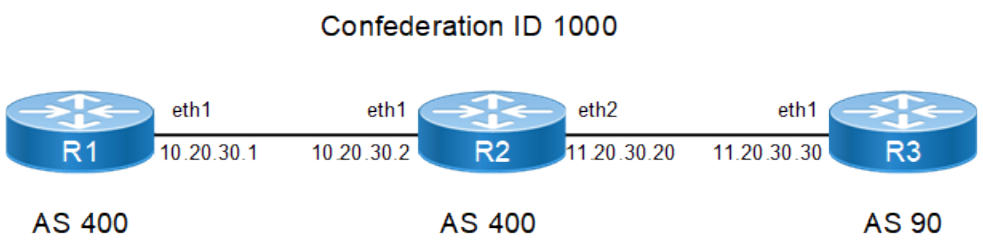

R1

#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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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

© 2026 IP Infusion Inc. Proprietary

`(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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

R2

The following provides the R2 validation:

```
#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

The following provides the R1 validation:

```
#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

The following provides the R3 validation:

```
#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 30. Single Confederation with Outside AS

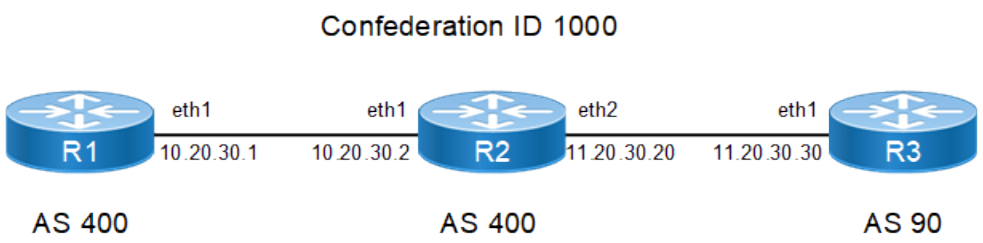

R1

`#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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

R2

`(config)#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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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

© 2026 IP Infusion Inc. Proprietary

`(config-router-af)#exit-address-family`	Exit address-family mode.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

R3

The following provides the R3 validation:

```
#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

The following provides the R2 validation:

```
#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.)
```

© 2026 IP Infusion Inc. Proprietary

R1

The following provides the R1 validation:

```
#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 31. IPv4 IBGP Peering

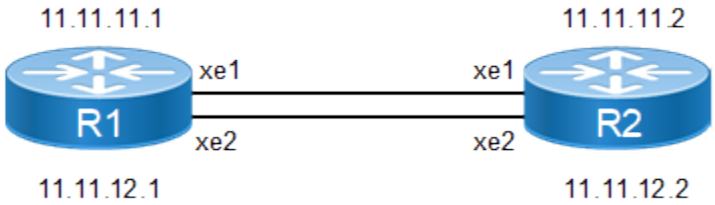

R1

`#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

© 2026 IP Infusion Inc. Proprietary

(config-router-af)#exit-address-family	Exit from address family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

R1

The following provides the R1 validation:

```
#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/(512)

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

Network   Next Hop   Metric   LocPrf   Weight Path
*>i   22.22.22.22/32   11.11.11.2   0   100   0   i
* i   11.11.12.2   0   100   0   i
*>   33.33.33.33/32   0.0.0.0   0   100   32768   i

Total number of prefixes 2
```

R2

The following provides the R2 validation:

```
#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 32. IPv4 VRF IBGP Peering

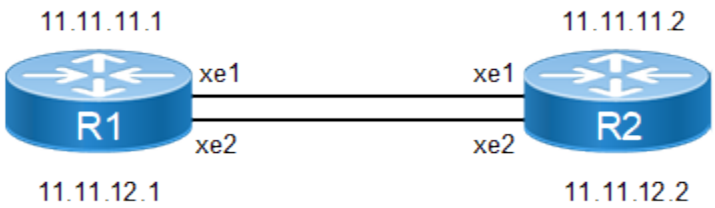

R1

#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)#exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

R2

#configure terminal (config)#ip vrf vrfA	Enter configure mode. 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.

© 2026 IP Infusion Inc. Proprietary

`#configure terminal` `(config)#interface xe2`	Enter configure mode. 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)#exit-address-family`	Exit address-family mode.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

R1

The following provides the R1 validation:

```
#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/(512)
```

```
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 33. 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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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.

© 2026 IP Infusion Inc. Proprietary

(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)#exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

Validation

R2

The following provides the R1 validation:

```
#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/(512)
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
 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 22.22.22.22/32
 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 34. IPv6 IBGP Peering

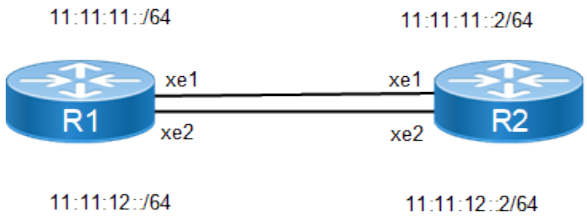

R1

`#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.
(config-router-af)#network 33::1/128	Advertise the loopback network into the BGP IPv6 address family.
(config-router-af)#exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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.

© 2026 IP Infusion Inc. Proprietary

`(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)#exit-address-family`	Exit address-family mode.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

R1

The following provides the R1 validation:

```
#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 35. IPv6 VRF IBGP peering

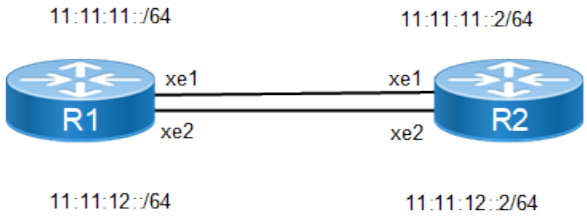

R1

`#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)#exit-address-family`	Exit address-family mode.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

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.
(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 11:11:12::1 activate	Activate neighbor in the address family
(config-router-af)#neighbor 11:11:11::1 remote-as 100	Configure BGP neighbor by specifying the neighbor IP address.
(config-router-af)# neighbor 11:11:11::1 activate	Activate neighbor in the address family
(config-router-af)#exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

R1

The following provides the R1 validation:

```
#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/(512)

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 36. 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)#exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

R2

#configure terminal	Enter Configuration 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.

© 2026 IP Infusion Inc. Proprietary

(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)#exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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)#exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

R2

The following provides the R1 validation:

```
#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/(512)
```

```
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
```

VPNV4 Configuration

Below mentioned topology displays bgp vpnv4 configuration on PE nodes, R1 and R3. IBGP peering will be formed on the loopback interface of R1 and R3; also IGP is running between all the routers.

Topology

Figure 37. IPv4 IBGP VPNv4 Configuration

R1

`#configure terminal`	Enter Configuration mode.
`(config)#ip vrf vrf1`	Create a VRF, vrf1.
`(config-vrf)#rd 100:1`	Configure a route distinguisher value.
`(config-vrf)#route-target export 100:1`	Configure a route target export value to VRF.
`(config-vrf)#route-target import 200:1`	Configure a route target import value to VRF.
`(config-vrf)#exit`	Exit from VRF configuration mode.
`(config)#router ldp`	Enter Router LDP mode.
`(config-router)#router-id 3.3.3.3`	Configure an LDP router ID.
`(config-router)#exit`	Exit from Router LDP mode.
`(config)#interface xe5`	Enter Interface configuration mode.

© 2026 IP Infusion Inc. Proprietary

`(config-if)#ip vrf forwarding vrf1`	Configure the interface to a VRF.
`(config-if)#ip address 1.1.1.1/24`	Assign an IP address to the interface.
`(config-if)#exit`	Exit from Interface configuration mode.
`(config-if)#interface xe1`	Enter another interface.
`(config-if)#ip address 11.11.11.1/24`	Assign an IP address to the interface.
`(config-if)#label-switching`	Enable label switching on interface.
`(config-if)#enable-ldp ipv4`	Enable IPv4 LDP configuration on interface.
`(config-if)#exit`	Exit from Interface configuration mode.
`(config-if)#interface lo`	Enter the loopback interface.
`(config-if)#ip address 20.20.20.20/32 secondary`	Assign a secondary IP address to the interface.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit from Interface Configuration mode.
`(config)#router ospf 100`	Enter Router OSPF mode.
`(config-router)#network 11.11.11.0/24 area 0`	Configure the interface on which OSPF runs, and associate the area ID.
`(config-router)#network 20.20.20.20/32 area 0`	Configure the interface on which OSPF runs, and associate the area ID.
`(config-router)#exit`	Exit from Router OSPF mode.
`(config)#router ospf 200 vrf1`	Create an OSPF process on VRF.
`(config-router)#network 1.1.1.1/24 area 0`	Configure the interface on which OSPF runs, and associate the area ID.
`(config-router)#redistribute bgp`	Redistribute BGP into OSPF.
`(config-router)#exit`	Exit from Router OSPF mode.
`(config)#router bgp 100`	Create a BGP process.
`(config-router)#neighbor lo_peer peer-group range 30.30.30.30/32`	Configure a dynamic peer group with the range command.
`(config-router)#neighbor lo_peer remote-as 100`	Configure remote AS to the peer group.
`(config-router)#neighbor lo_peer update-source lo`	Configure BGP neighbors to update the source routes.
`(config-router)#address-family vpnv4 unicast`	Enter the VPNv4 Address Family.
`(config-router-af)#neighbor lo_peer activate`	Activate the peer group in VPNv4 address family.
`(config-router-af)#exit-address-family`	Exit from VPNv4 address family.
`(config-router)#address-family ipv4 vrf vrf1`	Enter IPv4 VRF address family.
`(config-router-af)#redistribute ospf 200`	Redistribute OSPF into the IPv4 VRF address family.
`(config-router-af)#exit-address-family`	Exit address-family mode.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

R2

#configure terminal	Enter Configuration mode.
(config)#router ldp	Enter Router LDP mode.
(config-router)#router-id 4.4.4.4	Configure an LDP router ID.
(config-router)#exit	Exit from Router LDP mode.
(config-if)#interface xe2	Enter Interface Configuration mode.
(config-if)#ip address 12.12.12.2/24	Assign an IP address to the interface.
(config-if)#label-switching	Enable label switching on the interface.
(config-if)#enable-ldp ipv4	Enable IPv4 LDP configuration on the interface.
(config-if)#exit	Exit from Interface configuration mode.
(config-if)#interface xe1	Enter another Interface.
(config-if)#ip address 11.11.11.2/24	Assign an IP address to the interface.
(config-if)#label-switching	Enable label switching on the interface.
(config-if)#enable-ldp ipv4	Enable IPv4 LDP configuration on the interface.
(config-if)#exit	Exit from Interface mode.
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router ospf 100	Create an OSPF process.
(config-router)#network 11.11.11.0/24 area 0	Define the interface on which OSPF runs, and associate the area ID
(config-router)#network 12.12.12.0/24 area 0	Define the interface on which OSPF runs, and associate the area ID
(config-router)#exit	Exit from Router BGP mode.
(config)#commit	Commit the candidate configuration to the running configuration.

R3

#configure terminal	Enter Configuration mode.
(config)#router ldp	Enter Router LDP mode.
(config-router)#router-id 5.5.5.5	Configure an LDP router ID.
(config-router)#exit	Exit from Router LDP mode.
(config)#ip vrf vrf2	Create a VRF, vrf2.
(config-vrf)#rd 200:1	Configure a route distinguisher value.
(config-vrf)#route-target export 200:1	Configure a route target export value to VRF.
(config-vrf)#route-target import 100:1	Configure a route target import value to VRF.
(config-vrf)#exit	Exit from VRF configuration mode.
(config)#interface xe1	Enter Interface configuration mode.

© 2026 IP Infusion Inc. Proprietary

`(config-if)#ip vrf forwarding vrf2`	Configure an interface to a VRF.
`(config-if)#ip address 2.2.2.3/24`	Assign an IP address to the interface.
`(config-if)#exit`	Exit from Interface configuration mode.
`(config-if)#interface xe2`	Enter another interface.
`(config-if)#ip address 12.12.12.3/24`	Assign an IP address to the interface.
`(config-if)#label-switching`	Enable label switching on interface.
`(config-if)#enable-ldp ipv4`	Enable IPv4 LDP configuration on the interface.
`(config-if)#exit`	Exit from Interface configuration mode.
`(config-if)#interface lo`	Enter loopback interface.
`(config-if)#ip address 30.30.30.30/32 se`	Assign a secondary IP address to the interface.
`(config-if)#exit`	Exit from Interface mode.
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router ospf 100`	Enter Router OSPF mode.
`(config-router)#network 12.12.12.0/24 area 0`	Define the interface on which OSPF runs, and associate the area ID
`(config-router)#network 30.30.30.30/32 area 0`	Define the interface on which OSPF runs, and associate the area ID
`(config-router)#exit`	Exit from Router OSPF mode.
`(config)#router ospf 200 vrf2`	Create an OSPF process on VRF.
`(config-router)#network 2.2.2.3/24 area 0`	Define the interface on which OSPF runs, and associate the area ID.
`(config-router)#redistribute bgp`	Redistribute BGP into OSPF.
`(config-router)#exit`	Exit from Router OSPF mode.
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router bgp 100`	Create a BGP process.
`(config-router)#neighbor 20.20.20.20 remote-as 100`	Configure BGP neighbor by specifying a neighbor IP address.
`(config-router)#neighbor 20.20.20.20 update-s lo`	Define the BGP neighbors to update the source routes.
`(config-router)#address-family vpnv4 unicast`	Enter VPNv4 Address Family.
`(config-router-af)#neighbor 20.20.20.20 activate`	Activate the neighbor in VPNv4 address family.
`(config-router-af)#exit-address-family`	Exit from VPNv4 address family.
`(config-router)#address-family ipv4 vrf vrf2`	Enter IPv4 VRF address family.
`(config-router-af)#redistribute ospf 200`	Redistribute OSPF into the IPv4 address family.
`(config-router-af)#exit-address-family`	Exit address-family mode.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

Validation

R1

The following provides the R1 validation:

```
#show running-config router bgp
router bgp 100
 neighbor lo_peer peer-group range 30.30.30.30/32
 neighbor lo_peer remote-as 100
 neighbor lo_peer update-source lo
 !
 address-family vpnv4 unicast
 neighbor lo_peer activate
 exit-address-family
 !
 address-family ipv4 vrf vrf1
 redistribute ospf 200
 exit-address-family
!

#show ip bgp vpnv4 all summary
BGP router identifier 192.168.52.3, 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
*30.30.30.30             4   100    4          4       2      0      0  00:00:37               1
* Dynamically created based on a listen range command

BGP dynamic peer-group: lo_peer
  listen range:  30.30.30.30/32
  Total number of dynamically created neighbors/limit: 1/(512)

Total number of dynamically created neighbors: 1
Total number of activated dynamic peer-groups for VPNv4 Unicast address-family: 1

Total number of neighbors 1

Total number of Established sessions 1


#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: 100:1 (Default for VRF vrf1)
*>   1.1.1.0/24       0.0.0.0              1         100           32768    ?
*>i  2.2.2.0/24       30.30.30.30          1         100           0        ?
 Announced routes count = 1
 Accepted routes count = 1
Route Distinguisher: 200:1
*>i  2.2.2.0/24       30.30.30.30          1         100           0        ?
 Announced routes count = 0
 Accepted routes count = 1


#show ip bgp vpnv4 all 1.1.1.0
Route Distinguisher: 100:1
  Local
    20.20.20.20 (metric 12) from 20.20.20.20 (192.178.50.2)
      Origin incomplete, metric 1, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:100:1 0:0  OSPF-Route-type:0.0.0.0 :3:0
```

© 2026 IP Infusion Inc. Proprietary

```
      Last update: Tue Apr 23 10:29:10 2019

Route Distinguisher: 200:1 (Default for VRF vrf2)
  Local
    20.20.20.20 from 20.20.20.20 (192.178.50.2)
      Origin incomplete, metric 1, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:100:1 0:0  OSPF-Route-type:0.0.0.0 :3:0

      Last update: Tue Apr 23 10:29:10 2019


#show ip bgp peer-group

BGP dynamic peer-group is lo_peer, IBGP, remote AS 100
  BGP dynamic peer-group lo_peer listen range group members:
  30.30.30.30/32
  BGP version 4
Minimum time between advertisement runs is 5 seconds
For address family: IPv4 Unicast
  Peer-group member:
  *30.30.30.30
  Index 1, Offset 0, Mask 0x2
  0 accepted prefixes, 0 announced prefixes
For address family: VPNv4 Unicast
  Peer-group member:
  *30.30.30.30
  Index 0, Offset 0, Mask 0x0
  1 accepted prefixes, 1 announced prefixes
```

R2

The following provides the R2 validation:

```
R2#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: 100:1 (Default for VRF vrf1)
*>   1.1.1.0/24       0.0.0.0              1         100           32768    ?
*>i  2.2.2.0/24       30.30.30.30          1         100           0        ?
 Announced routes count = 1
 Accepted routes count = 1
Route Distinguisher: 200:1
*>i  2.2.2.0/24       30.30.30.30          1         100           0        ?
 Announced routes count = 0
 Accepted routes count = 1
R2#
```

R3

The following provides the R3 validation:

```
R3#show ip bgp vpnv4 all 1.1.1.0
Route Distinguisher: 100:1
  Local
    20.20.20.20 (metric 12) from 20.20.20.20 (192.178.50.2)
      Origin incomplete, metric 1, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:100:1 0:0  OSPF-Route-type:0.0.0.0 :3:0

      Last update: Tue Apr 23 10:29:10 2019
```

```
Route Distinguisher: 200:1 (Default for VRF vrf2)
  Local
    20.20.20.20 from 20.20.20.20 (192.178.50.2)
      Origin incomplete, metric 1, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:100:1 0:0  OSPF-Route-type:0.0.0.0 :3:0

      Last update: Tue Apr 23 10:29:10 2019

R3#
```

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 38. eBGP Multihop Connection

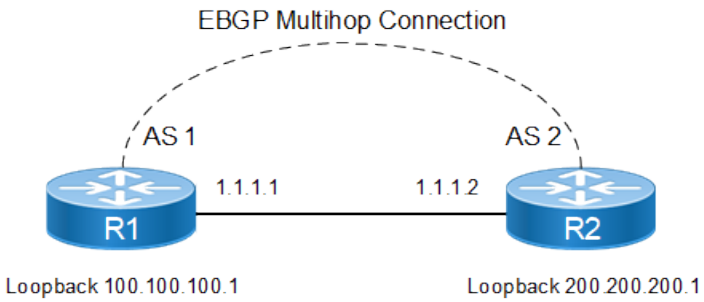

R1

`#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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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.
(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
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

R1

The following provides the R1 validation:

```
#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

The following provides the R2 validation:

```
#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 (page 1338)
- neighbor WORD as-origination-interval <1-65535> - neighbor as-origination-interval (page 1342)
- neighbor WORD attribute-unchanged ({as-path|next-hop|med}|) - neighbor attribute-unchanged (page 1343)
- neighbor WORD fall-over bfd - neighbor extended-optional-param (page 1364)
- neighbor WORD fall-over bfd multihop - neighbor extended-optional-param (page 1364)
- neighbor WORD next-hop-self - neighbor next-hop-self (page 1375)
- neighbor WORD remove-private-AS - neighbor remove-private-AS (page 1389)
- neighbor WORD route-reflector-client - neighbor route-reflector-client (page 1393)
- neighbor WORD route-server-client - neighbor route-server-client (page 1395)
- neighbor WORD send-community - neighbor send-community (page 1396)
- neighbor WORD distribute-list WORD out - neighbor distribute-list (page 1359)
- neighbor WORD capability orf prefix-list (both|receive|send) - neighbor capability orf prefix-list (page 1348)
- neighbor WORD filter-list WORD out - neighbor filter-list (page 1366)
- neighbor WORD prefix-list WORD out - neighbor prefix-list (page 1385)
- neighbor WORD route-map WORD out - neighbor route-map (page 1391)
- neighbor WORD advertisement-interval <1-65535> - neighbor advertisement-interval (page 1339)
- neighbor WORD disallow-infinite-holdtime - neighbor disallow-infinite-holdtime (page 1358)
- neighbor WORD local-as <1-4294967295> - neighbor local-as (page 1371)

Category B: 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 (page 1345)
- neighbor WORD remote-as <1-4294967295> - neighbor remote-as (page 1387)
- neighbor WORD allowas-in <1-10> - neighbor allowas-in (page 1340)
- neighbor WORD description WORD - neighbor description (page 1356)
- neighbor WORD distribute-list WORD In - neighbor distribute-list (page 1359)
- neighbor WORD ebgp-multihop - neighbor ebgp-multihop (page 1361)
- neighbor WORD ebgp-multihop <1-255> - neighbor ebgp-multihop (page 1361)
- neighbor WORD maximum-prefix <1-4294967295> - neighbor maximum-prefix (page 1373)

© 2026 IP Infusion Inc. Proprietary

- neighbor WORD update-source WORD - neighbor update-source (page 1406)
- neighbor WORD weight <0-65535> - neighbor weight (page 1409)
- neighbor WORD shutdown - neighbor shutdown (page 1399)
- neighbor WORD strict-capability-match - neighbor strict-capability-match (page 1401)
- neighbor WORD route-map WORD in - neighbor route-map (page 1391)
- neighbor WORD prefix-list WORD in - neighbor prefix-list (page 1385)
- neighbor WORD passive - neighbor passive (page 1378)
- neighbor WORD override-capability - neighbor override-capability (page 1377)
- neighbor WORD filter-list WORD in - neighbor filter-list (page 1366)
- neighbor WORD enforce-multihop - neighbor enforce-multihop (page 1363)
- neighbor WORD collide-established - neighbor collide-established (page 1350)

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 39. BGP Peer Groups with IPv4 Unicast Members

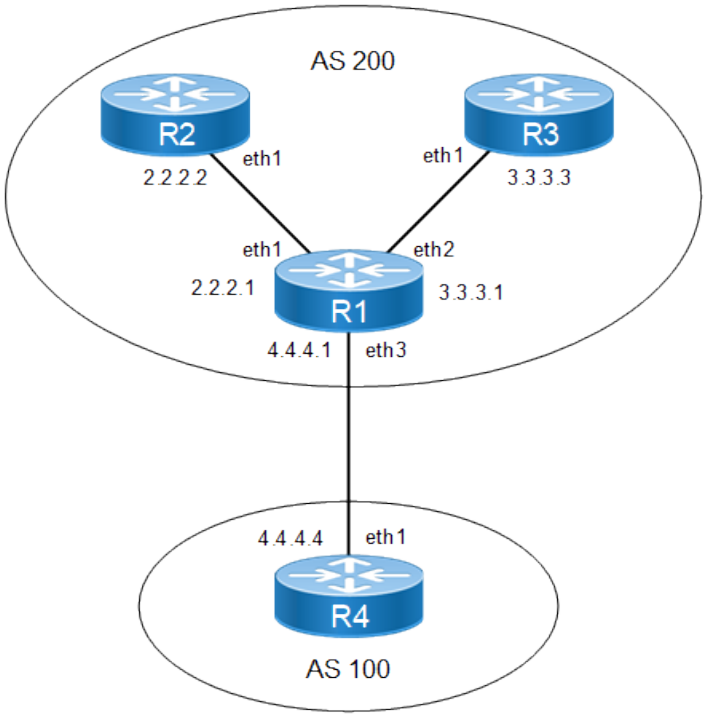

R1

`#configure terminal`	Enter configure mode.
`(config)#router bgp 200`	Define the routing process. The number 200 specifies the AS number of R2.

(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
(config-router-af)#exit-address-family	Exit address family mode
(config-router)#commit	Commit the candidate configuration to the running configuration.

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
(config-router-af)#exit-address-family	Exit address family mode
(config-router)#commit	Commit the candidate configuration to the running configuration.

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

© 2026 IP Infusion Inc. Proprietary

(config-router-af)#exit-address-family	Exit address family mode
(config-router)#commit	Commit the candidate configuration to the running configuration.

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
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

R1

The following provides the R1 validation:

```
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

The following provides the R2 validation:

```
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

The following provides the R3 validation:

```
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

© 2026 IP Infusion Inc. Proprietary

`(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)#commit`	Commit the candidate configuration to the running configuration.
`(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

The following provides the R1 validation:

```
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

The following provides the R2 validation:

```
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

The following provides the R3 validation:

```
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 address family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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-address-family	Exit address family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

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)#commit	Commit the candidate configuration to the running configuration.
(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

The following provides the R1 validation:

```
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-router)#commit	Commit the candidate configuration to the running configuration.
(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

The following provides the R1 validation:

```
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 40. Redistribute with OSPF

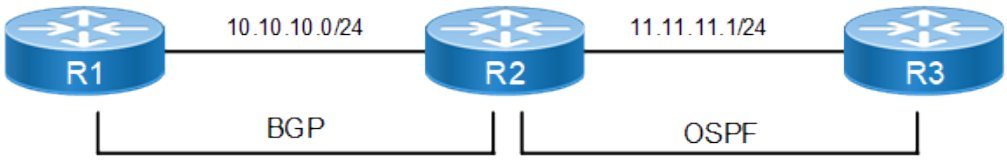

R2

`#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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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 41. Multiple Instances of Same AS

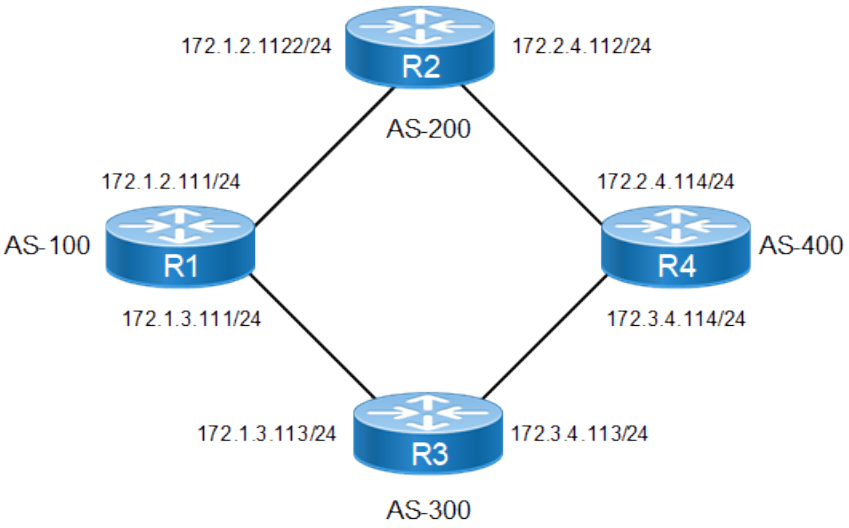

R1

`#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.
`(config-router-af)# exit-address-family`	Exit address-family mode.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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
(config-router-af)# exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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
(config-router-af)# exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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

© 2026 IP Infusion Inc. Proprietary

(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
(config-router-af)# exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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 42. 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

© 2026 IP Infusion Inc. Proprietary

(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-router)#commit	Commit the candidate configuration to the running configuration.
(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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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
(config-router)#commit	Commit the candidate configuration to the running configuration.

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
(config-router)#commit	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

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
(config-router)#commit	Commit the candidate configuration to the running configuration.

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
(config-router)#commit	Commit the candidate configuration to the running configuration.

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.

Notes:

- 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.
- OcNOS does not support 4-byte BGP Autonomous System Numbers (ASNs) in as-dot notation format

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 BGP MPLS VPN or 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.

Notes:

- Four-octet capability is disabled by default.
- OcNOS only supports 4-byte BGP ASNs in asplain format; as-dot notation is not supported.

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 43. 4-Octet ASN on Both Routers

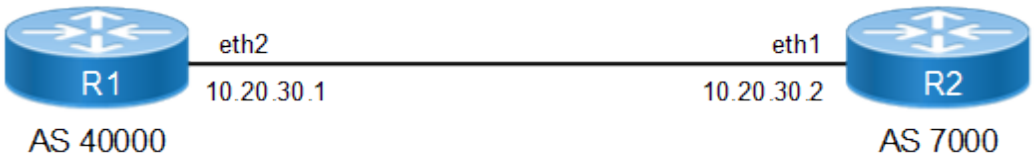

R1

`#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
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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 44. 4-Octet ASN on One Router

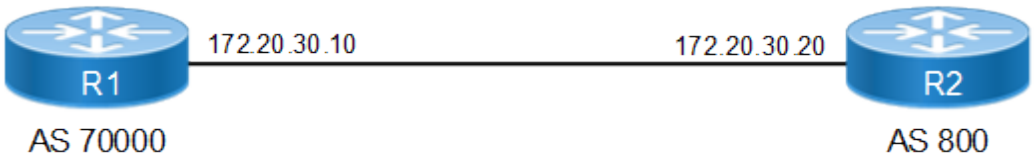

R1

`#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
(config-router)#commit	Commit the candidate configuration to the running configuration.

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
(config-router)#commit	Commit the candidate configuration to the running configuration.

Topology

Figure 45. 4-Octet ASN

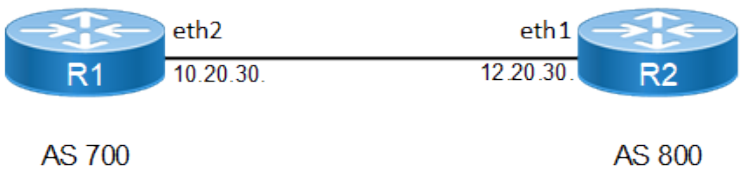

R1

#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

© 2026 IP Infusion Inc. Proprietary

(config-router-af)# exit-address-family	Exit address family mode
(config-router)#commit	Commit the candidate configuration to the running configuration.

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
(config-router)#commit	Commit the candidate configuration to the running configuration.

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
```

BGP Extended Community Attribute

The Extended Community Attribute provides a mechanism for labeling information carried in BGP.

Extended Community with a 2-Byte ASN

In the following example, CE1, PE1, PE2, and CE2 are 2-byte-ASN capable, and do not support 4-byte-ASN capability.

Topology

Figure 46. Extended Communities — 2-Byte ASN

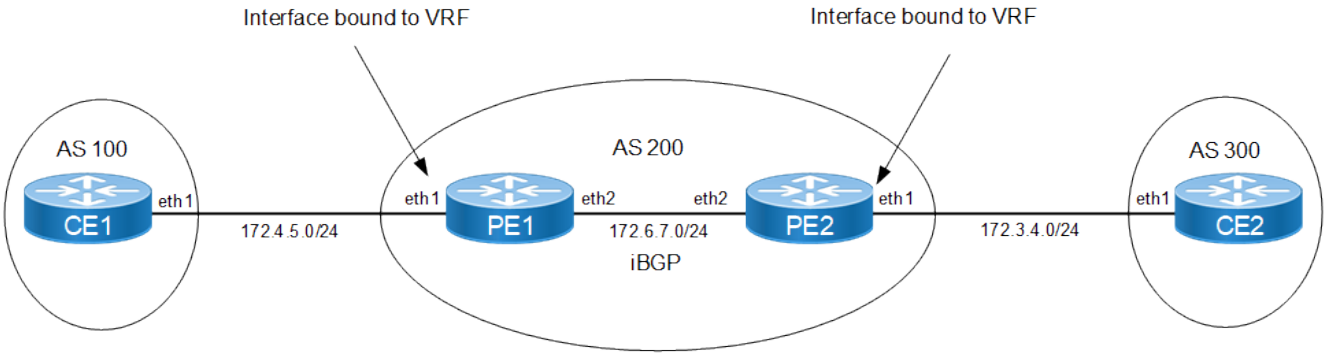

CE1

`#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)#router bgp 100`	Assign the ASN value (100) to the router. The ASN range is <1-65535>.
`(config-router)#neighbor 172.4.5.116 remote-as 200`	Specify the neighbor's IP address (172.4.5.116) and the ASN value of the neighbor (200).
`(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
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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)#router bgp 300`	Assign the ASN value (300) to the router. The ASN range is <1-65535>.
`(config-router)#neighbor 172.3.4.117 remote-as 200`	Specify the neighbor's IP address (172.3.4.117) and the ASN value of the neighbor (200).
`(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
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

PE1

`#configure terminal`	Enter configure mode.
`(config)#ip vrf VRF1`	Specify the name of the VRF (VRF1) to be created.
`(config-vrf)#rd 100:10`	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 100:10`	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 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 75.1.1.0/24 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 200`	Assign the ASN value (200) to the router.
`(config-router)#neighbor 172.6.7.117 remote-as 200`	Specify the neighbor’s (PE2) IP address (172.6.7.117) and the ASN value of the neighbor (200). 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.

© 2026 IP Infusion Inc. Proprietary

(config-router-af)#neighbor 172.6.7.117 activate	Activate the neighbor in address-family.
(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 100	Specify the neighbor's (CE1) IP address and ASN value.
(config-router-af)#neighbor 172.4.5.115 activate	Activate the neighbor in address-family
(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)# exit-address-family	Exit address family mode
(config-router)#commit	Commit the candidate configuration to the running configuration.

PE2

#configure terminal	Enter configure mode.
(config)#ip vrf VRF1	Specify the name of the VRF (VRF1) tobe created.
(config-vrf)#rd 100:10	Assign a route distinguisher (RD) for the VRF.
(config-vrf)#route-target both 100:10	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 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 100.1.1.0/24 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 200	Assign the ASN value (200) to the router.
(config-router)#neighbor 172.6.7.116 remote-as 200	Specify the neighbor's (PE1) IP address (172.6.7.116) and the ASN value of the neighbor (200). 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 300	Specify the neighbor's (CE2) IP address and ASN value.
(config-router-af)# neighbor 17.3.4.114 activate	Activate the neighbor in 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)# exit-address-family	Exit address family mode
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

CE1

The following provides the CE1 validation:

```
#show running-config
!
no service password-encryption
!
logging monitor 7
!
ip vrf management


!
ip domain-lookup feature telnet feature ssh
snmp-server enable snmp
snmp-server view all .1 included feature ntp
ntp enable
username ocnos role network-admin password encrypted $1$AUeGhbf0$HCHhxemCQ39LPYOjC.Kb7/ feature
rsyslog
!
interface lo
ip address 127.0.0.1/8 ipv6 address ::1/128 mtu 65536
!
interface eth0
ip address 192.168.52.2/24
!
interface eth1
ip address 172.4.5.115/24
!
interface eth2 shutdown
!
interface eth3 shutdown
!
interface eth4 shutdown
!
interface eth5 shutdown
!
```

```
router bgp 100
neighbor 172.4.5.116 remote-as 200
!
 address-family ipv4 unicast
 neighbor 172.4.5.116 activate
 exit-address-family
!
line con 0 login
line vty 0 39 login
!
end

#
#show ip bgp
BGP table version is 8, 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
*>   75.1.1.0/24   172.4.5.116   0   100   0   200   ?
*>   100.1.1.0/24   172.4.5.116   0   100   0   200   ?

Total number of prefixes 2 #
#show ip bgp neighbors
BGP neighbor is 172.4.5.116, remote AS 200, local AS 100, external link
BGP version 4, local router ID 192.168.52.2, remote router ID 172.4.5.116 BGP state =
Established, up for 00:04:22
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 131 messages, 1 notifications, 0 in queue
Sent 129 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 8, neighbor version 8 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: 172.4.5.115, Local port: 179
Foreign host: 172.4.5.116, Foreign port: 37982
Nexthop: 172.4.5.115 Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:04:54, due to BGP Notification received Notification Error Message: (Cease/Peer
Unconfigured.)


#show ip bgp vrf all
BGP table version is 8, 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
*>   75.1.1.0/24   172.4.5.116   0   100   0   200 ?
*>   100.1.1.0/24   172.4.5.116   0   100   0   200 ?

Total number of prefixes 2 #


#show ip bgp summary vrf all
BGP router identifier 192.168.52.2, local AS number 100 BGP table version is 8
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   200   168   165   8   0   0   00:22:04
2

Total number of neighbors 1

Total number of Established sessions 1
```

PE1

The following provides the PE1 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: 100:10 (Default for VRF VRF1)
*>   75.1.1.0/24      0.0.0.0             0         100           32768    ?
*>i  100.1.1.0/24     172.6.7.117          0         100           0        ?
 Announced routes count = 1
 Accepted routes count = 1
Route Distinguisher: 100:10
*>i  100.1.1.0/24     172.6.7.117          0         100           0        ?
 Announced routes count = 0
 Accepted routes count = 1


#show ip bgp vrf all
BGP table version is 2, local router ID is 172.4.5.116
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
*>   75.1.1.0/24      0.0.0.0             0         100           32768    ?
*>i  100.1.1.0/24     172.6.7.117          0         100           0        ?

Total number of prefixes 2


#show ip bgp summary vrf all
BGP router identifier 172.4.5.116, local AS number 200
BGP VRF VRF1 Route Distinguisher: 100:10
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
172.4.5.115              4   100   55         60       2      0      0  00:26:54               0

Total number of neighbors 1

Total number of Established sessions 1
BGP router identifier 192.168.52.3, 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
```

```
172.6.7.117              4   200   80        101       1      0      0  00:37:47              0

Total number of neighbors 1

Total number of Established sessions 1


#show ip bgp neighbors
BGP neighbor is 172.6.7.117, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 192.168.52.3, remote router ID 192.168.52.5
  BGP state = Established, up for 00:38:33
  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
    Address family VPNv4 Unicast: advertised and received
  Received 82 messages, 0 notifications, 0 in queue
  Sent 103 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

 For address family: VPNv4 Unicast
  BGP table version 4, neighbor version 4
  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: 172.6.7.116, Local port: 179
Foreign host: 172.6.7.117, Foreign port: 57743
Nexthop: 172.6.7.116
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 172.4.5.115, vrf VRF1, remote AS 100, local AS 200, external link
  BGP version 4, local router ID 172.4.5.116, remote router ID 192.168.52.2
  BGP state = Established, up for 00:27:40
  Last read 00:00:10, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 57 messages, 0 notifications, 0 in queue
  Sent 62 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)
  0 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 172.4.5.116, Local port: 37982
Foreign host: 172.4.5.115, Foreign port: 179
Nexthop: 172.4.5.116
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
```

```
#show ip bgp vrf all
BGP table version is 2, local router ID is 172.4.5.116
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
*>   75.1.1.0/24      0.0.0.0             0         100           32768    ?
*>i  100.1.1.0/24     172.6.7.117          0         100           0        ?

Total number of prefixes 2
```

PE2

The following provides the PE2 validation:

```
#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  75.1.1.0/24      172.6.7.116          0         100           0        ?
*>   100.1.1.0/24     0.0.0.0             0         100           32768    ?

Total number of prefixes 2


#show ip bgp summary vrf all
BGP router identifier 172.3.4.117, local AS number 200
BGP VRF VRF1 Route Distinguisher: 100:10
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
172.3.4.114             4   300   82         85       1      0      0  00:40:05              0

Total number of neighbors 1

Total number of Established sessions 1
BGP router identifier 192.168.52.5, 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
172.6.7.116             4   200  113        113       1      0      0  00:54:07              0

Total number of neighbors 1

Total number of Established sessions 1#


#show ip bgp neighbors
BGP neighbor is 172.6.7.116, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 192.168.52.5, remote router ID 192.168.52.3
  BGP state = Established, up for 00:56:09
  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
    Address family VPNv4 Unicast: advertised and received
  Received 117 messages, 0 notifications, 0 in queue
  Sent 117 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
```

CE2

The following provides the CE2 validation:

```
#show ip bgp vpnv4 all
#show ip bgp
BGP table version is 3, 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
*>   75.1.1.0/24      172.3.4.117           0         100           0       200 ?
*>   100.1.1.0/24     172.3.4.117           0         100           0       200 ?

Total number of prefixes 2
#
#
#show ip bgp vrf all
BGP table version is 3, 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
*>   75.1.1.0/24      172.3.4.117           0         100           0       200 ?
*>   100.1.1.0/24     172.3.4.117           0         100           0       200 ?

Total number of prefixes 2
#
#
#show ip bgp summary vrf all
BGP router identifier 192.168.52.4, local AS number 300
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
172.3.4.117              4   200  382        414       3      0      0  00:42:54               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 200, local AS 300, external link
  BGP version 4, local router ID 192.168.52.4, remote router ID 172.3.4.117
  BGP state = Established, up for 00:43:04
  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 236 messages, 147 notifications, 0 in queue
  Sent 415 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: 172.3.4.114, Local port: 179
Foreign host: 172.3.4.117, Foreign port: 54753
Nexthop: 172.3.4.114
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:43:32, due to BGP Notification received
Notification Error Message: (OPEN Message Error/Bad Peer AS.)
#
```

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 47. Extended Communities — 4-Byte ASN

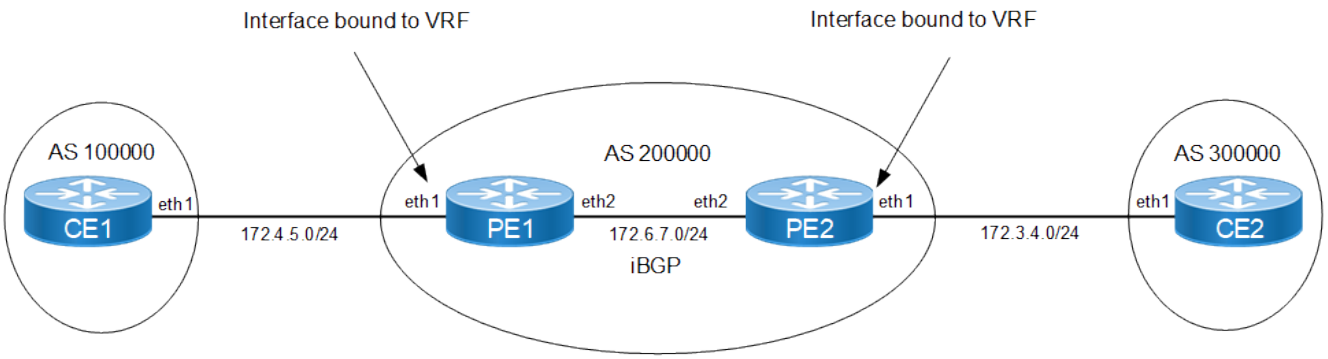

CE1

`#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
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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)# exit-address-family	Exit address family mode
(config-router)#commit	Commit the candidate configuration to the running configuration.

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)#commit	Commit the candidate configuration to the running configuration.
(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

© 2026 IP Infusion Inc. Proprietary

	(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)# exit-address-family`	Exit address family mode
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

Validation

CE1

The following provides the CE1 validation:

```
#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

The following provides the PE1 validation:

```
#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

The following provides the PE2 validation:

```
#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

The following provides the CE2 validation:

```
#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 48. BGP Nexthop Tracking

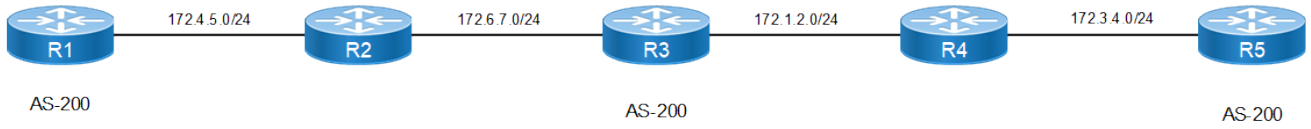

R1

`#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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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.

© 2026 IP Infusion Inc. Proprietary

(config-router)#commit	Commit the candidate configuration to the running configuration.

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.
(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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

R4

#configure terminal	Enter configure mode.
(config)#router ospf 1	Configure the OSPF process (1).

© 2026 IP Infusion Inc. Proprietary

(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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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)#commit	Commit the candidate configuration to the running configuration.
(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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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 49. Topology for Nexthop Tracking Delay Timer

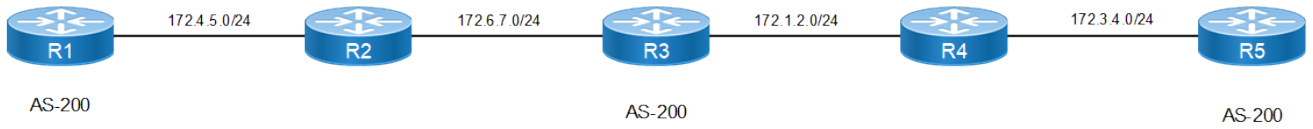

R1

`#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)#commit`	Commit the candidate configuration to the running configuration.
`(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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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-router)#commit`	Commit the candidate configuration to the running configuration.
`(config)#bgp nexthop-trigger enable`	Enable nexthop tracking.
`(config)#bgp nexthop-trigger delay 20`	Configure the nexthop trigger-delay time interval.

© 2026 IP Infusion Inc. Proprietary

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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

R1

The following provides the R1 validation:

```
#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

The following provides the R3 validation:

```
#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

The following provides the R5 validation:

```
#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 50. Administrative Distance for IPv4 BGP

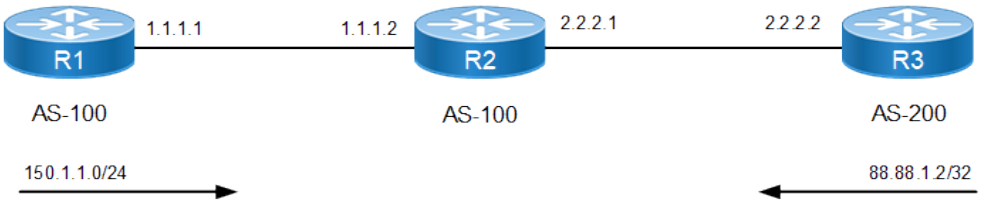

R1

`#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
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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.

© 2026 IP Infusion Inc. Proprietary

(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
(config-router)#commit	Commit the candidate configuration to the running configuration.

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
(config-router)#commit	Commit the candidate configuration to the running configuration.

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 51. BGP Weight Per Peer

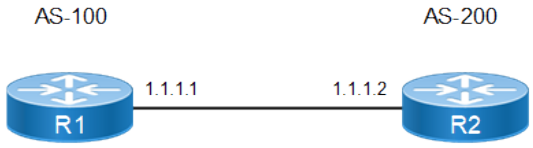

R1

`#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
(config-router)#commit	Commit the candidate configuration to the running configuration.

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
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

R1

The following provides the R1 validation:

```
#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

The following provides the R2 validation:

```
#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 MPLS VPN environment, customer networks are connected to an MPLS VPN-enabled provider backbone. As shown in Figure 52, Customer A areas, Areas 0 and 1, are connected to an MPLS 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 52. OSPF as PE-CE Protocol

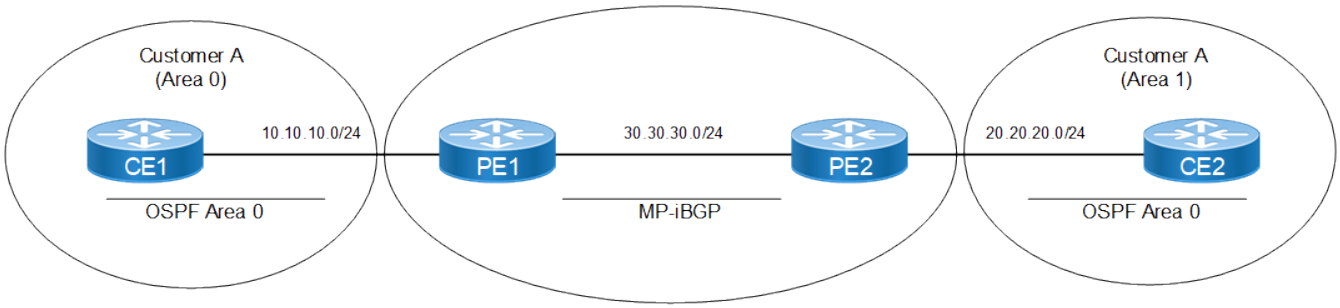

Configuration

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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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.

© 2026 IP Infusion Inc. Proprietary

(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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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)#commit	Commit the candidate configuration to the running configuration.
(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)#commit	Commit the candidate configuration to the running configuration.
(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.

© 2026 IP Infusion Inc. Proprietary

`(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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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 53. Multipath iBGP for IPv4

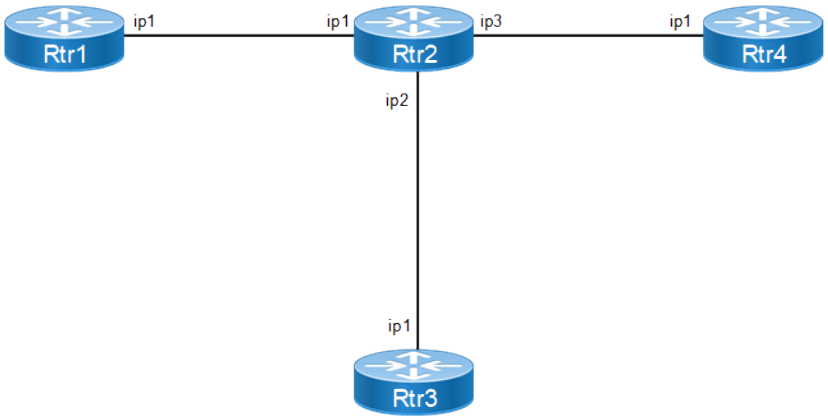

Configuration

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.
`(config)#commit`	Commit the candidate configuration to the running configuration.

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

© 2026 IP Infusion Inc. Proprietary

(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.
(config)#commit	Commit the candidate configuration to the running configuration.

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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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
(config-router)#commit	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

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 Mar   2 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 Mar   2 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 Mar   2 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
neighbor 30.30.30.2   remote-as   100
neighbor 40.40.40.4   remote-as   100
neighbor 50.50.50.6   remote-as   100

!
 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 54. Multipath eBGP for IPv4

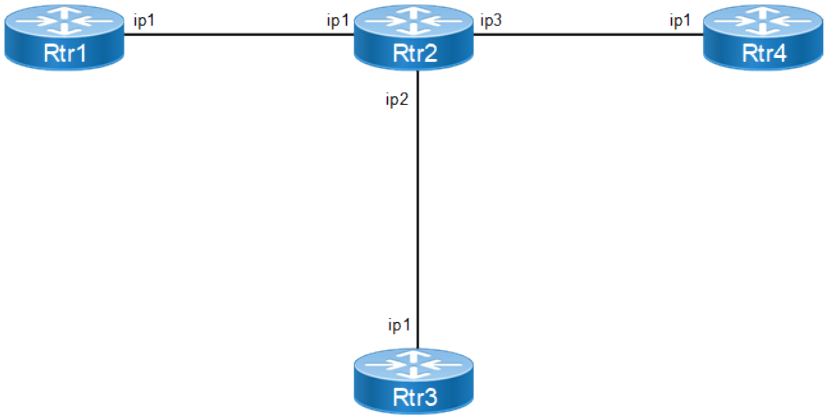

Configuration

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.
`(config)#commit`	Commit the candidate configuration to the running configuration.

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

© 2026 IP Infusion Inc. Proprietary

(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.
(config)#commit	Commit the candidate configuration to the running configuration.

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.
(config)#commit	Commit the candidate configuration to the running configuration.

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
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

© 2026 IP Infusion Inc. Proprietary

`(config-router)#exit`	Exit the Router mode and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration.

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 Jan   3 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
Jan   3 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 Jan   3 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 55. Multipath eiBGP topology

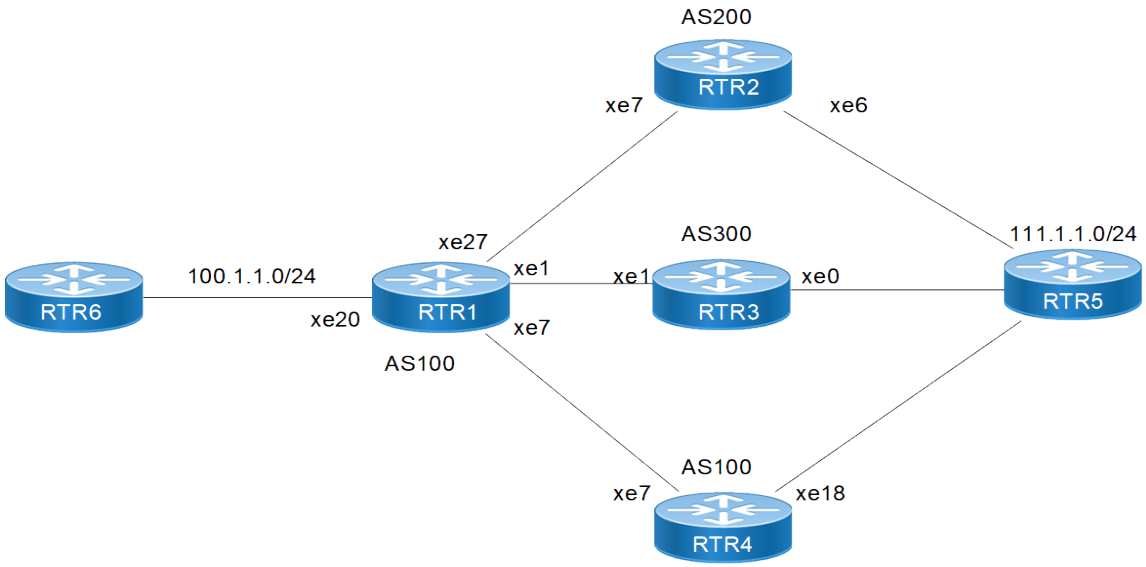

Configuration

RTR1

`#configure terminal`	Enter Configure mode.
`(config)#interface lo`	Enter Interface mode
`(config-if)# ip address 41.41.41.41/32 sec-ondary`	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 acti-vate`	Activate the neighbor
`(config-router-af)#neighbor 22.1.1.2 acti-vate`	Activate the neighbor
`(config-router-af)#neighbor 23.1.1.2 acti-vate`	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 sec-ondary`	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

© 2026 IP Infusion Inc. Proprietary

`(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 acti-vate`	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 sec-ondary`	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 acti-vate`	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

© 2026 IP Infusion Inc. Proprietary

(config-if)# ip address 43.43.43.43/32 sec-ondary	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 acti-vate	Activate the neighbor
(config-router-af)# commit	Commit the configurations
(config-router-af)# end	Return to privilege mode

Validation

RTR1

The following provides the RTR1 validation:

```
# 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
```

BGP 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

Figure 56. BGP AS-PATH Multipath-relax Topology

Below topology explains about BGP AS PATH multipath relax functionality.

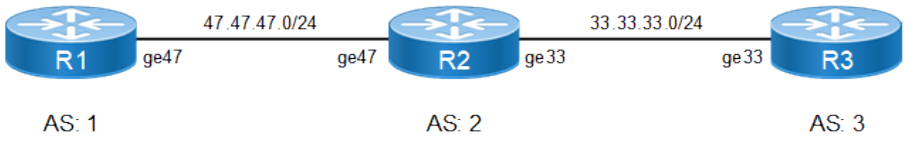

Configuration

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)#exit-address-family	Exit address-family config mode
(config-router)#commit	Commit the candidate configuration to the running configuration.

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

© 2026 IP Infusion Inc. Proprietary

	configure mode
(config-router)# bgp bestpath as-path multipath-relax	Configure BGP AS PATH Multipath relax.
(config-router-af)#exit-address-family	Exit address-family mode
(config-router)#commit	Commit the candidate configuration to the running configuration.

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)#exit-address-family	Exit from router BGP and address-family config mode
(config-router)#commit	Commit the candidate configuration to the running configuration.

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

Neighbor   V   AS   MsgRcv   MsgSen TblVer   InQ   OutQ   Up/ Down   State/PfxRcd
33.33.33.3
00:01:10   1
47.47.47.1
00:06:33   1

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

Network   Next Hop   Metric   LocPrf   Weight   Path
*>   100.1.1.0/24   47.47.47.1   0   100   0      1 i
*   33.33.33.3   0   100   0      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.

© 2026 IP Infusion Inc. Proprietary

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 57. BFP FIB Install Topology

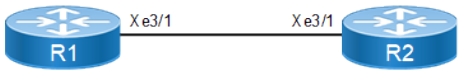

Configuration

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.

© 2026 IP Infusion Inc. Proprietary

(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.
(config)#commit	Commit the candidate configuration to the running configuration.

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)#exit-address-family	Exit address family mode
(config-router)#commit	Commit the candidate configuration to the running configuration.

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
Neighbor   V   AS   MsgRcv   MsgSen TblVer   InQ   OutQ   Up/Dow
n   State/PfxRcd
20.1.1.2   4   100   5
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


   Network   Next Hop   Metric   LocPrf   Weight Path
*>   1.1.1.0/24   0.0.0.0   0   100   32768   ?
*>   2.2.2.0/24   0.0.0.0   0   100   32768   ?
*>   3.3.3.0/24   0.0.0.0   0   100   32768   ?
*>   4.4.4.0/24   0.0.0.0   0   100   32768   ?
*>   5.5.5.0/24   0.0.0.0   0   100   32768   ?
*>   6.6.6.0/24   0.0.0.0   0   100   32768   ?

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]   is   directly   connected,   eth1,   00:06:54
S   2.2.2.0/24   [1/0]   is   directly   connected,   eth1,   00:06:35
S   3.3.3.0/24   [1/0]   is   directly   connected,   eth1,   00:06:26
S   4.4.4.0/24   [1/0]   is   directly   connected,   eth1,   00:06:17
S   5.5.5.0/24   [1/0]   is   directly   connected,   eth1,   00:06:09
S   6.6.6.0/24   [1/0]   is   directly   connected,   eth1,   00:06:01
C   20.1.1.0/24 is directly connected, eth1, 00:07:32 C   127.0.0.0/8 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

   Network   Next Hop   Metric   LocPrf   Weight Path
*>   1.1.1.0/24   0.0.0.0   0   100   32768   ?
*>   2.2.2.0/24   0.0.0.0   0   100   32768   ?
*>   3.3.3.0/24   0.0.0.0   0   100   32768   ?
*>   4.4.4.0/24   0.0.0.0   0   100   32768   ?
```

```
*>   5.5.5.0/24   0.0.0.0   0   100   32768   ?
*>   6.6.6.0/24   0.0.0.0   0   100   32768   ?


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]   via   20.1.1.1,   eth1,   00:13:44
B   2.2.2.0/24   [200/0]   via   20.1.1.1,   eth1,   00:13:44
B   3.3.3.0/24   [200/0]   via   20.1.1.1,   eth1,   00:13:44
B   4.4.4.0/24   [200/0]   via   20.1.1.1,   eth1,   00:13:44
B   5.5.5.0/24   [200/0]   via   20.1.1.1,   eth1,   00:13:44
B   6.6.6.0/24   [200/0]   via   20.1.1.1,   eth1,   00:13:44
C   20.1.1.0/24 is directly connected, eth1, 00:14:12 C   127.0.0.0/8 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

   Network   Next Hop   Metric   LocPrf   Weight   Path
*>i   1.1.1.0/24   20.1.1.1   0   100   0      ?
*>i   2.2.2.0/24   20.1.1.1   0   100   0      ?
*>i   3.3.3.0/24   20.1.1.1   0   100   0      ?
*>i   4.4.4.0/24   20.1.1.1   0   100   0      ?


*>i   5.5.5.0/24   20.1.1.1   0   100   0   ?
*>i   6.6.6.0/24   20.1.1.1   0   100   0   ?

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

Neighbor   V   AS   MsgRcv   MsgSen TblVer   InQ   OutQ   Up/Down   State/ PfxRcd
20.1.1.1   4   100   40   39   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 C   127.0.0.0/8 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)#exit-address-family`	Exit address family mode
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

```
#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]   via   20.1.1.1,   eth1,   00:00:04
B   2.2.2.0/24   [200/0]   via   20.1.1.1,   eth1,   00:00:04
B   3.3.3.0/24   [200/0]   via   20.1.1.1,   eth1,   00:00:04
B   4.4.4.0/24   [200/0]   via   20.1.1.1,   eth1,   00:00:04
B   5.5.5.0/24   [200/0]   via   20.1.1.1,   eth1,   00:00:04
B   6.6.6.0/24   [200/0]   via   20.1.1.1,   eth1,   00:00:04
C   20.1.1.0/24 is directly connected, eth1, 00:31:16 C   127.0.0.0/8 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 and this feature is not supported for VPNV4 address-family

Route Target Constraint

BGP/MPLS IP VPNs use PE routers to Route Target (RT) extended communities and control the distribution of routes into the VRFs. Within a given iBGP mesh, PE routers hold routes marked with RouteTargets pertaining to VRFs that have local CE attachments.

BGP RT Constrained Route Distribution is a feature that can be used by service providers in Multiprotocol Label Switching (MPLS) Layer 3 VPNs to reduce the number of unnecessary routing updates that route reflectors (RRs) send to Provider Edge (PE) routers. The reduction in “routing updates” saves resources by allowing RRs, Autonomous System Boundary Routers (ASBRs), and PEs to carry fewer routes. Route targets are used to constrain routing updates.

With (MPLS)VPNs, the (iBGP) peers or Route Reflectors send all VPN4 and/or VPN6 prefixes to the PE routers. The PE routers drop the VPN4/6 prefixes for which there is no importing VPN route forwarding (VRF).

Topology

The topology below shows Route-target filtering in an L3VPN—with Route Target Constraint (RTC), the RR sends only wanted VPN4/6 prefixes to the PE; wanted” means that the PEs have the VRFs importing the specific prefixes.

Figure 58. Route-target Filter Topology

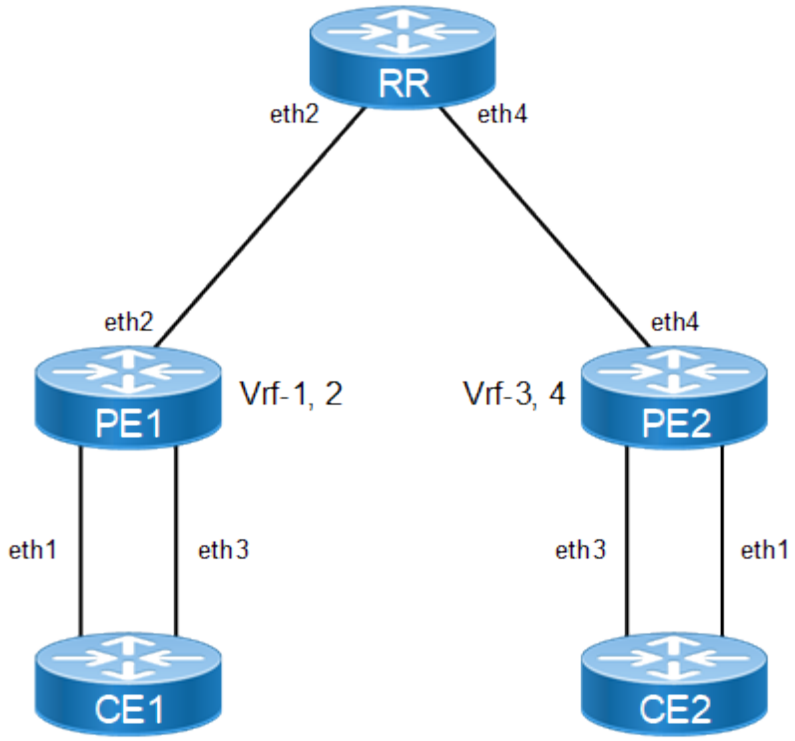

Configuration

CE1

#configure terminal	Enter configure mode.
(config)#interface eth1	Enter interface mode
(config-if)#ip address 80.1.1.1/24	Configure IP address for interface
(config-if)#exit	Exit from interface mode and enter into Configure mode
(config)#interface eth3	Enter interface mode
(config-if)#ip address 90.1.1.1/24	Configure IP address for interface
(config-if)#exit	Exit from interface mode and enter into Configure mode
(config)# router bgp 200	Assign the ASN value (100) to the BGP router
(config-router)#neighbor 80.1.1.2 remote-as 100	Configure neighbor (RR) in IBGP
(config-router)#neighbor 90.1.1.2 remote-as 100	Configure neighbor (RR) in IBGP
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)#redistribute static	Redistribute static routes into BGP
(config-router-af)#neighbor 80.1.1.2 activate	Activate neighbor
(config-router-af)#neighbor 90.1.1.2 activate	Activate neighbor
(config-router-af)#exit-address-family	Exit address-family mode
(config-router)#exit	Exit from router mode and enter configure mode
(config)#ip route vrf 1 1.1.1.0/24 eth1	Configure static route with VRF 1 instance
(config)#ip route vrf 2 3.3.3.0/24 eth3	Configure static route with VRF 2 instance
(config)#ip route vrf 2 4.4.4.0/24 eth3	Configure static route with VRF 2 instance
(config)#commit	Commit the candidate configuration to the running configuration.

CE2

#configure terminal	Enter configure mode.
(config)#interface eth1	Enter interface mode
(config-if)#ip address 101.1.1.1/24	Configure IP address for interface
(config-if)#exit	Exit from interface mode and enter into Configure mode
(config)#interface eth3	Enter interface mode
(config-if)#ip address 100.1.1.1/24	Configure IP address for interface
(config-if)#exit	Exit from interface mode and enter into Configure mode

© 2026 IP Infusion Inc. Proprietary

(config)# router bgp 200	Assign the ASN value (100) to the BGP router
(config-router)#neighbor 100.1.1.2 remote-as 100	Configure neighbor (RR) in IBGP
(config-router)#neighbor 101.1.1.2 remote-as 100	Configure neighbor (RR) in IBGP
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)#neighbor 100.1.1.2 activate	Activate neighbor
(config-router-af)#neighbor 101.1.1.2 activate	Activate neighbor
(config-router-af)#exit-address-family	Exit from address family ipv4 unicast mode
(config-router)#exit	Exit from router and configure mode
(config)#commit	Commit the candidate configuration to the running configuration.

PE1

#configure terminal	Enter configure mode.
(config)#ip vrf 1	Create a VRF instance 1
(config-vrf)#rd 1:100	Configure unique RD value for VRF to identify VRF instance
(config-vrf)#route-target export 1:200	Configure route-target (rt) value for exporting routes into other VRFs (for other PE's)
(config-vrf)#exit	Exit VRF mode and enter Configure mode
(config)#ip vrf 2	Create a VRF instance 2
(config-vrf)#rd 1:300	Configure unique RD value for VRF to identify VRF instance
(config-vrf)#route-target both 1:400	Configure route-target (rt) value for exporting routes into other VRFs (for other PE's)
(config-vrf)#exit	Exit VRF mode and enter Configure mode
(config)#router ldp	Enable LDP.
(config-router)#exit	Exit router LDP mode
(config)#interface lo	Enter loopback interface mode
(config-if)#ip address 11.11.11.11/32 secondary	Configure IP address for loopback interface
(config-if)# enable-ldp ipv4	Enable LDP on loopback interface
(config-if)#exit	Exit interface mode
(config)#interface eth1	Enter interface mode
(config-if)#ip vrf forwarding 1	Bind interface to VRF 1
(config-if)#ip address 80.1.1.2/24	Configure IP address for VRF binded interface
(config-if)#exit	Exit interface mode
(config)#interface eth2	Enter interface mode

© 2026 IP Infusion Inc. Proprietary

(config-if)#ip address 40.1.1.1/24	Configure an IP address for interface
(config-if)#label-switching	Enable label-switching on interface
(config-if)# enable-ldp ipv4	Enable LDP on connected interface between PE1 and RR
(config-if)#exit	Exit interface mode
(config)#interface eth3	Enter interface mode
(config-if)#ip vrf forwarding 2	Bind interface to VRF 1
(config-if)#ip address 90.1.1.2/24	Configure IP address for VRF binded interface
(config-if)#exit	Exit interface mode
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router ospf	Enable OSPF process between PE1 and RR
(config-router)#network 11.11.11.11/32 area 0.0.0.0 (config-router)#network 40.1.1.0/24 area 0.0.0.0	Advertise loopback network in OSPF area 0
(config-router)#exit	Exit router OSPF mode
(config)#commit	Commit the candidate configuration to the running configuration.
(config)# router bgp 100	Assign the ASN value (100) to the BGP router
(config-router)#neighbor 22.22.22.22 remote-as 100	Configure neighbor (RR) in IBGP
(config-router)#neighbor 22.22.22.22 update-source lo	Enable neighbor with loopback interface.
(config-router)#address-family vpnv4 unicast	Enter Address-Family-VPNv4 mode.
(config-router-af)#neighbor 22.22.22.22 activate	Activate RR neighbor
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#address-family rtfilter unicast	Enable RT filter address-family mode
(config-router-af)#neighbor 22.22.22.22 activate	Activate neighbor
(config-router-af)#exit-address-family	Exit RTfilter Address Family mode and return to Router mode.
(config-router)#address-family ipv4 vrf 1	Enter Address-Family-VRF mode.
(config-router-af)#neighbor 80.1.1.1 remote-as 200	Configure CE neighbor in VRF mode
(config-router-af)#neighbor 80.1.1.1 activate	Activate neighbor in VRF
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#address-family ipv4 vrf 2	Enter Address-Family-VRF mode.
(config-router-af)#neighbor 90.1.1.1 remote-as 200	Configure CE neighbor in VRF mode
(config-router-af)#neighbor 90.1.1.1 activate	Activate neighbor in VRF
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router

© 2026 IP Infusion Inc. Proprietary

	mode.
(config-router)#exit	Exit from router mode and configure mode
(config-router)#commit	Commit the candidate configuration to the running configuration.

RR

(config)#router ldp	Enable LDP
(config-router)#exit	Exit router LDP mode
(config)#interface lo	Enter loopback interface
(config-if)#ip address 22.22.22.22/32 secondary (config-if)#ip address 44.44.44.44/32 secondary	Configure IP address for loopback interface
(config-if)# enable-ldp ipv4	Enable LDP on loopback interface
(config-if)#exit	Exit interface mode
(config)#interface eth2	Enter interface mode
(config-if)#ip address 40.1.1.2/24	Configure IP address for interface connecting to PE2
(config-if)#label-switching	Enable label-switching on interface
(config-if)# enable-ldp ipv4	Enable LDP on connected interface between PE1 and RR
(config-if)#exit	Exit interface mode
(config)#interface eth4	Enter into interface mode
(config-if)#ip address 50.1.1.1/24	Configure an IP address for interface connecting to PE1
(config-if)#label-switching	Enable label-switching on interface
(config-if)# enable-ldp ipv4	Enable LDP on connected interface between PE1 and RR
(config-if)#exit	Exit interface mode
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router ospf	Enable OSPF process between PE1 and RR
(config-router)#network 22.22.22.22/32 area 0.0.0.0	Advertise loopback network in OSPF area 0
(config-router)#network 40.1.1.0/24 area 0 (config-router)#network 44.44.44.44/32 area 0.0.0.0 (config-router)#network 50.1.1.0/24 area 0.0.0.0	Advertise PE1 to RR connected network in OSPF
(config-router)#exit	Exit from router OSPF mode
(config)#commit	Commit the candidate configuration to the running configuration.
(config)# router bgp 100	Assign the ASN value (100) to the BGP router

(config-router)#neighbor 11.11.11.11 remote-as 100	Configure neighbor (PE1) in IBGP
(config-router)#neighbor 11.11.11.11 update-source 22.22.22.22	Enable neighbor with loopback interface
(config-router)#neighbor 33.33.33.33 remote-as 100	Configure neighbor (PE2) in IBGP
(config-router)#neighbor 33.33.33.33 update-source 44.44.44.44	Enable neighbor with loopback interface
(config-router)#address-family vpnv4 unicast	Enter Address-Family-VPNv4 mode.
(config-router-af)#neighbor 11.11.11.11 activate	Activate PE1 neighbor
(config-router-af)#neighbor 33.33.33.33 activate	Activate PE2 neighbor
(config-router-af)#neighbor 11.11.11.11 route-reflector-client	Configure PE1 as Route Reflector client
(config-router-af)#neighbor 33.33.33.33 route-reflector-client	Configure PE2 as Route Reflector client
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#address-family rtfilter unicast	Enable RT filter address-family mode
(config-router-af)#neighbor 11.11.11.11 activate	Activate PE1 neighbor in RTfilter family
(config-router-af)#neighbor 33.33.33.33 activate	Activate PE2 neighbor in RTfilter family
(config-router-af)#neighbor 33.33.33.33 route-reflector-client	Configure PE2 as Route Reflector client
(config-router-af)#neighbor 11.11.11.11 route-reflector-client	Configure PE1 as Route Reflector client
(config-router-af)#exit-address-family	Exit RTfilter Address-Family mode
(config-router)#exit	Exit from Address-Family, Router and Configure mode.
(config)#commit	Commit the candidate configuration to the running configuration.

PE2

#configure terminal	Enter configure mode.
(config)#ip vrf 3	Create a VRF instance 3
(config-vrf)#rd 1:600	Configure unique RD value for VRF to identify VRF instance
(config-vrf)#route-target export 1:200	Configure route-target (rt) value for exporting routes into other VRFs (for other PE's)
(config-vrf)#exit	Exit VRF mode and enter Configure mode
(config)#ip vrf 4	Create a VRF instance 4
(config-vrf)#rd 1:900	Configure unique RD value for VRF to identify VRF instance
(config-vrf)#route-target both 1:400	Configure route-target (rt) value for exporting routes into other VRFs (for other PE's)

© 2026 IP Infusion Inc. Proprietary

`(config-vrf)#exit`	Exit VRF mode and enter Configure mode
`(config)#router ldp`	Enable LDP.
`(config-router)#exit`	Exit router LDP mode
`(config)#interface lo`	Enter loopback interface mode
`(config-if)#ip address 33.33.33.33/32 secondary`	Configure IP address for loopback interface
`(config-if)# enable-ldp ipv4`	Enable LDP on loopback interface
`(config-if)#exit`	Exit interface mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip vrf forwarding 3`	Bind interface to VRF 3
`(config-if)#ip address 101.1.1.2/24`	Configure IP address for VRF binded interface
`(config-if)#exit`	Exit interface mode
`(config)#interface eth3`	Enter interface mode
`(config-if)#ip vrf forwarding 4`	Bind interface to VRF 3
`(config-if)#ip address 100.1.1.2/24`	Configure IP address for VRF binded interface
`(config-if)#exit`	Exit interface mode
`(config)#interface eth4`	Enter interface mode
`(config-if)#ip address 50.1.1.2/24`	Configure an IP address for interface
`(config-if)#label-switching`	Enable label-switching on interface
`(config-if)# enable-ldp ipv4`	Enable LDP on connected interface between PE2 and RR
`(config-if)#exit`	Exit interface mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router ospf`	Enable OSPF process between PE2 and RR
`(config-router)#network 33.33.33.33/32 area 0.0.0.0`	Advertise loopback network in OSPF area 0
`(config-router)#network 50.1.1.0/24 area 0`	Advertise PE2 to RR connected network in OSPF
`(config-router)#exit`	Exit router OSPF mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)# router bgp 100`	Assign the ASN value (100) to the BGP router
`(config-router)#neighbor 44.44.44.44 remote-as 100`	Configure neighbor (RR) in IBGP
`(config-router)#neighbor 44.44.44.44 update-source 33.33.33.33`	Enable neighbor with loopback interface.
`(config-router)#address-family vpnv4 unicast`	Enter Address-Family-VPNv4 mode.
`(config-router-af)#neighbor 44.44.44.44 activate`	Activate RR neighbor
`(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router

© 2026 IP Infusion Inc. Proprietary

	mode.
(config-router)#address-family rtfilter unicast	Enable RT filter address-family mode
(config-router-af)#neighbor 44.44.44.44 activate	Activate neighbor
(config-router-af)#exit-address-family	Exit RTfilter Address Family mode and return to Router mode.
(config-router)#address-family ipv4 vrf 3	Enter Address-Family-VRF mode.
(config-router-af)#neighbor 101.1.1.1 remote-as 200	Configure CE neighbor in VRF mode
(config-router-af)#neighbor 101.1.1.1 activate	Activate neighbor in VRF
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#address-family ipv4 vrf 4	Enter Address-Family-VRF mode.
(config-router-af)#neighbor 100.1.1.1 remote-as 200	Configure CE neighbor in VRF mode
(config-router-af)#neighbor 100.1.1.1 activate	Activate neighbor in VRF
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#exit	Exit router and configure mode
(config)#commit	Commit the candidate configuration to the running configuration.

Validation

Through RTfilter, address-family RT values will be exchanged between RR and PEs. Neighbors are activated under this address-family and configured clients as well. RR will learn routes from PEs and send them to other PEs if it has any peer requesting those routes based on their RT import values.

The outputs below show the routes sent and learned in PEs, installed in VRFs, and display the RT filter values exchanged between them. For more details on show command output fields, refer to show ip bgp (page 1560), show ip bgp vpnv4 (page 1573), and show ip bgp rtfilter all (page 1567)

CE1

The following provides the CE1 validation:

```
CE1#show ip bgp
BGP table version is 6, local router ID is 192.160.50.5
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       0.0.0.0              0         100          32768    ?
*>   3.3.3.0/24       0.0.0.0              0         100          32768    ?
*>   4.4.4.0/24       0.0.0.0              0         100          32768    ?

Total number of prefixes 3
```

PE1

The following provides the PE1 validation:

```
PE1#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: 1:100 (Default for VRF 1)
*>   1.1.1.0/24       80.1.1.1              0         100           0       200 ?
*>   3.3.3.0/24       80.1.1.1              0         100           0       200 ?
*>   4.4.4.0/24       80.1.1.1              0         100           0       200 ?
 Announced routes count = 3
 Accepted routes count = 0
Route Distinguisher: 1:300 (Default for VRF 2)
*>   1.1.1.0/24       90.1.1.1              0         100           0       200 ?
*>   3.3.3.0/24       90.1.1.1              0         100           0       200 ?
*>   4.4.4.0/24       90.1.1.1              0         100           0       200 ?
 Announced routes count = 3
 Accepted routes count = 0

PE1#show ip bgp rtfilter all
RTFilter's Received
*******************
peer-ip 22.22.22.22
100:1:400/96

RTFilter's Sent
*******************
peer-ip 22.22.22.22
100:1:400/96

RTFilter's rt_add_pending
*******************

RTFilter's rt_del_pending
*******************

RTFilter's rtfilter_receive_pending
*******************

 Example: The format 100:1:400/96 is interpreted as ORIGIN_AS:Route-Target/Prefix-length
 Breakdown of 100:1:400/96:
        The first part (100) is the Global Administrator ASN represents Origin AS.
        1:400 -> Represents the Route Target (RT) Extended Community.
        /96 -> The prefix length, which in the case of RTC is always 96 bits.
```

RR

The following provides the RR validation:

```
RR#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: 1:300
*>i  1.1.1.0/24       11.11.11.11           0         100           0       200 ?
*>i  3.3.3.0/24       11.11.11.11           0         100           0       200 ?
*>i  4.4.4.0/24       11.11.11.11           0         100           0       200 ?
 Announced routes count = 0
 Accepted routes count = 3
```

```
RR#show ip bgp rtfilter all
RTFilter's Received
*******************
peer-ip 11.11.11.11
100:2:1:400/96
peer-ip 33.33.33.33
100:2:1:400/96

RTFilter's Sent
*******************
peer-ip 11.11.11.11
100:2:1:400/96
peer-ip 33.33.33.33
100:2:1:400/96

RTFilter's rt_add_pending
*******************

RTFilter's rt_del_pending
*******************

RTFilter's rtfilter_receive_pending
*******************

 Example: The format 100:1:400/96 is interpreted as ORIGIN_AS:Route-Target/Prefix-length
 Breakdown of 100:1:400/96:
         The first part (100) is the Global Administrator ASN represents Origin AS.
         1:400 -> Represents the Route Target (RT) Extended Community.
         /96 -> The prefix length, which in the case of RTC is always 96 bits.
```

PE2

The following provides the PE2 validation:

```
PE2#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: 1:300
*>i  1.1.1.0/24       11.11.11.11          0         100           0       200 ?
*>i  3.3.3.0/24       11.11.11.11          0         100           0       200 ?
*>i  4.4.4.0/24       11.11.11.11          0         100           0       200 ?
 Announced routes count = 0
 Accepted routes count = 3
Route Distinguisher: 1:900 (Default for VRF 4)
*>i  1.1.1.0/24       11.11.11.11          0         100           0       200 ?
*>i  3.3.3.0/24       11.11.11.11          0         100           0       200 ?
*>i  4.4.4.0/24       11.11.11.11          0         100           0       200 ?
 Announced routes count = 0
 Accepted routes count = 3

PE2#show ip bgp rtfilter all
RTFilter's Received
*******************
peer-ip 44.44.44.44
100:2:1:400/96
RTFilter's Sent
*******************
peer-ip 44.44.44.44
100:2:1:400/96

RTFilter's rt_add_pending
```

```
******************

RTFilter's rt_del_pending
******************

RTFilter's rtfilter_receive_pending
******************

 Example: The format 100:1:400/96 is interpreted as ORIGIN_AS:Route-Target/Prefix-length
 Breakdown of 100:1:400/96:
         The first part (100) is the Global Administrator ASN represents Origin AS.
         1:400 -> Represents the Route Target (RT) Extended Community.
         /96 -> The prefix length, which in the case of RTC is always 96 bits.
```

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.

Notes: 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.

Notes: Beware of these 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).

Notes: Beware of these items:

- 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.

© 2026 IP Infusion Inc. Proprietary

Notes: Beware of these items:

- 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.

Notes: 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`

Notes: 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.

Notes: Beware of these items:

- 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 59. BGP dampening

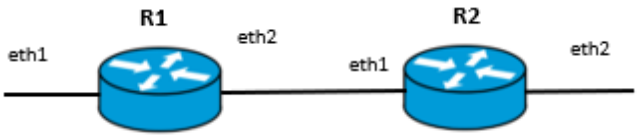

IPv4 Configuration

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)#commit`	Commit the candidate configuration to the running configuration.
`(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

© 2026 IP Infusion Inc. Proprietary

(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
(config-router)#commit	Commit the candidate configuration to the running configuration.

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)#commit	Commit the candidate configuration to the running configuration.
(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
`(config-router-af)#exit-address-family`	Exit address-family mode
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

R2

The following provides the R2 validation:

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

   Network   From   Reuse   Path
d   200.1.0.0   10.1.1.1   00:29:00 100   300   i
d   200.2.0.0   10.1.1.1   00:28:20 100   300   i
d   200.3.0.0   10.1.1.1   00:28:20 100   300   i
d   200.4.0.0   10.1.1.1   00:28:20 100   300   i
d   200.5.0.0   10.1.1.1   00:28:20 100   300   i
d   200.6.0.0   10.1.1.1   00:28:20 100   300   i
d   200.7.0.0   10.1.1.1   00:28:20 100   300   i
d   200.8.0.0   10.1.1.1   00:28:20 100   300   i
d   200.9.0.0   10.1.1.1   00:28:20 100   300   i
d   200.10.0.0   10.1.1.1   00:28:20 100   300   i
```

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

   Network   From   Flaps   Duration   Reuse   Path
d   200.1.0.0   10.1.1.1   8   00:18:37   00:29:10   100 300   i
d   200.2.0.0   10.1.1.1   7   00:14:22   00:29:00   100 300   i
d   200.3.0.0   10.1.1.1   7   00:14:22   00:29:00   100 300   i
d   200.4.0.0   10.1.1.1   7   00:14:22   00:29:00   100 300   i
```

```
d   200.5.0.0   10.1.1.1   7   00:14:22   00:29:00   100 300   i
d   200.6.0.0   10.1.1.1   7   00:14:22   00:29:00   100 300   i
d   200.7.0.0   10.1.1.1   7   00:14:22   00:29:00   100 300   i
d   200.8.0.0   10.1.1.1   7   00:14:22   00:29:00   100 300   i
d   200.9.0.0   10.1.1.1   7   00:14:22   00:29:00   100 300   i
d   200.10.0.0   10.1.1.1   7   00:14:22   00:29:00   100 300   i
```

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
`(config-router-af)#exit-address-family`	Exit address-family mode
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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

© 2026 IP Infusion Inc. Proprietary

`(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
`(config-router-af)#exit-address-family`	Exit address-family mode
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

R2

The following provides the R2 validation:

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 60. BGP Authentication

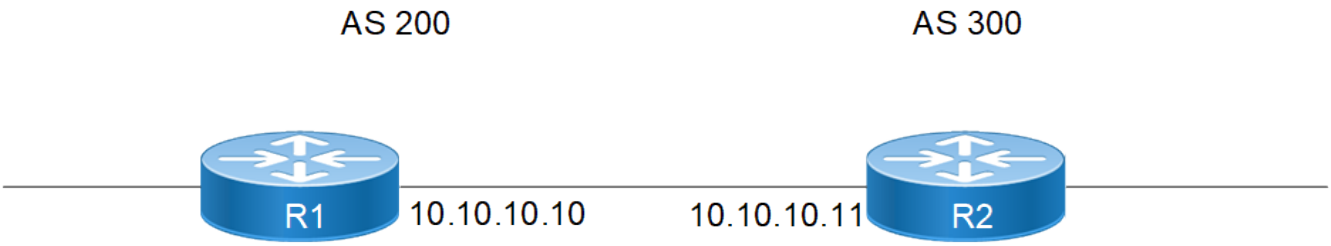

Configuration

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)#exit`	End config mode
`(config)#commit`	Commit the candidate configuration to the running configuration.

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

© 2026 IP Infusion Inc. Proprietary

(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)#exit	End config mode
(config)#commit	Commit the candidate configuration to the running configuration.

Validation

R1

The following provides the R1 validation:

```
#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

The following provides the R2 validation:

```
#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

Th is 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.

© 2026 IP Infusion Inc. Proprietary

Notes:

- 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.
- 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.
- 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.
- 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 61. BGP unnumbered

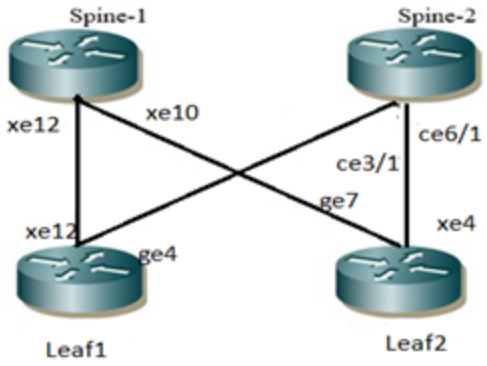

Configuration

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
(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

© 2026 IP Infusion Inc. Proprietary

(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

© 2026 IP Infusion Inc. Proprietary

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:

The following provides the Spine1 validation:

```
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:

The following provides the Spine2 validation:

```
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
```

© 2026 IP Infusion Inc. Proprietary

```
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

The following provides the Leaf 1 validation:

```
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

The following provides the Leaf 2 validation:

```
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.

Notes:

- 65535:666 is reserved for Blackhole community.
- BGP blackhole community is supported only for unicast address-family.

Topology

Figure 62. BGP Blackhole Community Attribute topology

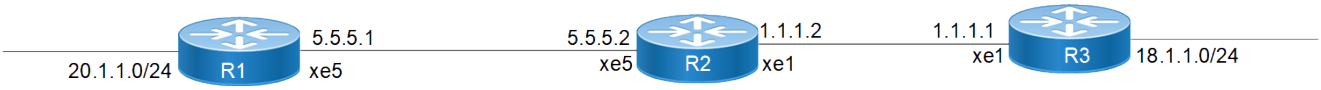

Configuration

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

© 2026 IP Infusion Inc. Proprietary

(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

© 2026 IP Infusion Inc. Proprietary

(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

The following provides the R2 validation:

```
# 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

The following provides the R1 validation:

```
#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
```

Accumulated Interior Gateway Protocol

Overview

This section includes step-by-step configuration of AIGP Metric Attribute for BGP It also contains an overview of the concepts of AIGP. AIGP is a non-transitive attribute that includes the accumulated IGP metric. BGP routers advertise this AIGP metric to neighbors in other ASes. This allows BGP routers to select the best path based on the end-to-end IGP metric.

Topology

The diagram depicts the topology for the configuration examples that follow.

Figure 63. Accumulated Interior Gateway Protocol

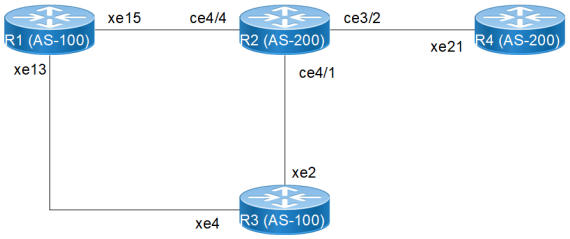

Base Configurations

R1

Loopback Interface:

`#configure terminal`	Enter configuration mode.
`(config)#interface lo`	Enter the Interface mode for the loopback interface.
`(config-if)#ip address 1.1.1.1/32 secondary`	Configure IP address on loopback interface.
`(config-if)#commit`	Commit the transaction

Interface Configuration:

`(config)#interface xe15`	Enter the Interface mode for xe15.
`(config-if)# ip address 10.10.10.1/24`	Configure IP address on the interface.
`(config-if)#exit`	Exit interface mode
`(config)#interface xe13`	Enter the Interface mode for xe13.
`(config-if)# ip address 11.11.11.1/24`	Configure IP address on the interface.
`(config)#interface xe12`	Enter the Interface mode for xe12.
`(config-if)# ip address 50.50.50.2/24`	Configure IP address on the interface.
`(config-if)#commit`	Commit the transaction.

OSPF Configuration:

`(config)#router ospf 100`	Enter the Router OSPF mode.
`(config-router)#ospf router-id 1.1.1.1`	Router-id configurations
`(config-router)# network 1.1.1.1/32 area 0`	Advertise loopback address in OSPF.
`(config-router)#network 10.10.10.0/24 area 0`	Advertise network address in OSPF.
`(config-router)#network 11.11.11.0/24 area 0`	Advertise network address in OSPF.
`(config-router)#commit`	Commit the transaction

© 2026 IP Infusion Inc. Proprietary

R2

Loopback Interface:

`#configure terminal`	Enter configuration mode.
`(config)#interface lo`	Enter the Interface mode for the loopback interface.
`(config-if)#ip address 2.2.2.2/32 secondary`	Configure IP address on loopback interface.
`(config-if)#commit`	Commit the transaction

Interface Configuration:

`(config)#interface ce4/4`	Enter the Interface mode for ce4/4.
`(config-if)# ip address 10.10.10.2/24`	Configure IP address on the interface.
`(config-if)#interface ce4/1`	Enter the Interface mode for ce4/1.
`(config-if)# ip address 12.12.12.1/24`	Configure IP address on the interface.
`(config-if)#interface ce3/2`	Enter the Interface mode for ce3/2.
`(config-if)# ip address 13.13.13.1/24`	Configure IP address on the interface.
`(config-if)#commit`	Commit the transaction.

OSPF Configuration:

`(config)#router ospf 100`	Enter the Router OSPF mode.
`(config-router)#ospf router-id 2.2.2.2`	Router-id configurations
`(config-router)# network 2.2.2.2/32 area 0`	Advertise loopback address in OSPF.
`(config-router)#network 10.10.10.0/24 area 0`	Advertise network address in OSPF.
`(config-router)#network 12.12.12.0/24 area 0`	Advertise network address in OSPF.
`(config-router)#network 13.13.13.0/24 area 0`	Advertise network address in OSPF.
`(config-router)#commit`	Commit the transaction

R3

Loopback Interface:

`#configure terminal`	Enter configuration mode.
`(config)#interface lo`	Enter the Interface mode for the loopback interface.
`(config-if)#ip address 3.3.3.3/32 secondary`	Configure IP address on loopback interface.
`(config-if)#commit`	Commit the transaction

Interface Configuration:

`(config)#interface xe4`	Enter the Interface mode for xe4.

(config-if)# ip address 11.11.11.2/24	Configure IP address on the interface.
(config)#interface xe2	Enter the Interface mode for xe2.
(config-if)# ip address 12.12.12.2/24	Configure IP address on the interface.
(config-if)#commit	Commit the transaction.

OSPF Configuration:

(config)#router ospf 100	Enter the Router OSPF mode.
(config-router)#ospf router-I 3.3.3.3	Router-id configurations
(config-router)# network 3.3.3.3/32 area 0	Advertise loopback address in OSPF.
(config-router)#network 11.11.11.0/24 area 0	Advertise network address in OSPF.
(config-router)#network 12.12.12.0/24 area 0	Advertise network address in OSPF.
(config-router)#commit	Commit the transaction

R4

Loopback Interface:

#configure terminal	Enter configuration mode.
(config)#interface lo	Enter the Interface mode for the loopback interface.
(config-if)#ip address 4.4.4.4/32 secondary	Configure IP address on loopback interface.
(config-if)#commit	Commit the transaction

Interface Configuration:

(config)#interface xe12	Enter the Interface mode for xe21.
(config-if)# ip address 13.13.13.2/24	Configure IP address on the interface.
(config-if)#commit	Commit the transaction.

OSPF Configuration:

(config)#router ospf 100	Enter the Router OSPF mode.
(config-router)#ospf router-id 4.4.4.4	Router-id configurations
(config-router)# network 4.4.4.4/32 area 0	Advertise loopback address in OSPF.
(config-router)#network 13.13.13.0/24 area 0	Advertise network address in OSPF.
(config-router)#commit	Commit the transaction

© 2026 IP Infusion Inc. Proprietary

BGP Configurations

R1

Route Map Configuration:

`(config)# route-map map1 permit 10`	Configure route-map map1 with sequence number 10.
`(config-route-map)#set aigp-metric 100`	Set aigp metric as 100.
`(config-route-map)#commit`	Commit the transaction

BGP Configuration:

`(config)# router bgp 100`	Enter the BGP configuration mode.
`(config-router)# neighbor 3.3.3.3 remote-as 100`	Configure neighbor
`(config-router)# neighbor 3.3.3.3 update-source 1.1.1.1`	Update loopback address as source
`(config-router)# address-family ipv4 unicast`	Enter address family mode.
`(config-router)#network 50.50.50.0/24 route-map map1`	Apply a route map to routes.
`(config-router-af)#neighbor 3.3.3.3 activate`	Activate the neighbor.
`(config-router-af)#exit`	Exit address family mode.
`(config-router)#exit`	Exit Router BGP mode

R2

BGP Configuration:

`(config)# router bgp 200`	Enter the BGP configuration mode.
`(config-router)#neighbor 3.3.3.3 remote-as 100`	Configure neighbor
`(config-router)#neighbor 4.4.4.4 remote-as 200`	Configure neighbor
`(config-router)# neighbor 3.3.3.3 update-source 2.2.2.2`	Update loopback address as source
`(config-router)# neighbor 4.4.4.4 update-source 2.2.2.2`	Update loopback address as source
`(config-router)# neighbor 3.3.3.3 ebgp-multihop 2`	Enable multihop on eBGP session.
`(config-router)#address-family ipv4`	Enter address family mode.
`(config-router-af)#neighbor 3.3.3.3 activate`	Activate the neighbor.
`(config-router-af)#neighbor 4.4.4.4 activate`	Activate the neighbor.
`(config-router-af)#neighbor 3.3.3.3 aigp enable`	Enable aigp for the neighbor.
`(config-router-af)#neighbor 4.4.4.4 aigp enable`	Enable aigp for the neighbor.
`(config-router-af)#neighbor 4.4.4.4 next-hop-self`	Configure next-hop-self for the R4 neighbor.
`(config-router-af)#commit`	Commit the transaction.

© 2026 IP Infusion Inc. Proprietary

R3

BGP Configuration:

`(config)# router bgp 100`	Enter the BGP configuration mode.
`(config-router)#neighbor 1.1.1.1 remote-as 100`	Configure neighbor
`(config-router)#neighbor 2.2.2.2 remote-as 200`	Configure neighbor
`(config-router)#neighbor 4.4.4.4 remote-as 200`	Configure neighbor
`(config-router)#neighbor 1.1.1.1 update-source 3.3.3.3`	Update loopback address as source
`(config-router)#neighbor 2.2.2.2 update-source 3.3.3.3`	Update loopback address as source
`(config-router)#neighbor 4.4.4.4 update-source 3.3.3.3`	Update loopback address as source
`(config-router)# neighbor 2.2.2.2 ebgp-multihop 2`	Enable multihop on eBGP session.
`(config-router)# neighbor 4.4.4.4 ebgp-multihop 2`	Enable multihop on eBGP session.
`(config-router)#address-family ipv4`	Enter address family mode.
`(config-router-af)#neighbor 1.1.1.1 activate`	Activate the neighbor.
`(config-router-af)#neighbor 2.2.2.2 activate`	Activate the neighbor.
`(config-router-af)#neighbor 4.4.4.4 activate`	Activate the neighbor.
`(config-router-af)#neighbor 2.2.2.2 aigp enable`	Enable aigp for the neighbor.
`(config-router-af)#neighbor 4.4.4.4 aigp enable`	Enable aigp for the neighbor.
`(config-router-af)#commit`	Exit address family mode.

R4

BGP Configuration:

`(config)# router bgp 200`	Enter the BGP configuration mode.
`(config-router)#neighbor 3.3.3.3 remote-as 100`	Configure neighbor
`(config-router)#neighbor 2.2.2.2 remote-as 200`	Configure neighbor
`(config-router)# neighbor 3.3.3.3 update-source 4.4.4.4`	Update loopback address as source
`(config-router)# neighbor 2.2.2.2 update-source 4.4.4.4`	Update loopback address as source
`(config-router)# neighbor 3.3.3.3 ebgp-multihop 2`	Enable multihop on eBGP session.
`(config-router)#address-family ipv4`	Enter address family mode.
`(config-router-af)#neighbor 3.3.3.3 activate`	Activate the neighbor.
`(config-router-af)#neighbor 2.2.2.2 activate`	Activate the neighbor.

© 2026 IP Infusion Inc. Proprietary

(config-router-af)#neighbor 3.3.3.3 aigp enable	Enable aigp for the neighbor.
(config-router-af)#neighbor 2.2.2.2 aigp enable	Enable aigp for the neighbor.
(config-router-af)#neighbor 2.2.2.2 next-hop-self	Configure next-hop-self for the R4 neighbor.
(config-router-af)#commit	Commit the transaction.

Validation

R3

```
R3#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
1.1.1.1           1   Full/DR         00:00:37    11.11.11.1      xe4                     0
2.2.2.2           1   Full/Backup     00:00:37    12.12.12.1      xe2                     0
R3#sh ip bgp summary
BGP router identifier 3.3.3.3, local AS number 100
BGP table version is 23
1 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  162        172      23      0      0  00:23:34              1
2.2.2.2                  4   200   60         68      23      0      0  00:08:50              0
4.4.4.4                  4   200  145        164      23      0      0  00:03:19              0

Total number of neighbors 3

Total number of Established sessions 3
R3#sh ip bgp
BGP table version is 23, local router ID is 3.3.3.3
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
*>i  50.50.50.0/24    1.1.1.1               0        100       0    i

Total number of prefixes 1
R3#sh ip bgp 50.50.50.0
BGP routing table entry for 50.50.50.0/24
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  2.2.2.2 4.4.4.4
  Local
   Nexthop:1.1.1.1 (metric 2) from 1.1.1.1 (Remote Id:1.1.1.1)
      Origin IGP, metric 0, localpref 100, aigp metric 100      valid, internal, best, source safi: 1
      Not advertised to any peer
      Total AIGP metric 102
      Last update: Wed Jun 22 06:41:17 2022
```

R4

```
R4#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
2.2.2.2           1   Full/Backup     00:00:37    13.13.13.1      xe21                    0
R4#sh ip bgp summary
BGP router identifier 4.4.4.4, local AS number 200
```

```
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
2.2.2.2                  4   200   82         85       4      0      0  00:04:28              1
3.3.3.3                  4   100  100         96       4      0      0  00:04:28              1

Total number of neighbors 2

Total number of Established sessions 2
R4#sh ip bgp
BGP table version is 4, local router ID is 4.4.4.4
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
*>   50.50.50.0/24    3.3.3.3              0        100       0   100 i
* i                   2.2.2.2              0        100       0   100 i

Total number of prefixes 1
R4#sh ip bgp 50.50.50.0
BGP routing table entry for 50.50.50.0/24
Paths: (2 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  2.2.2.2
  AS path:100
   Nexthop:3.3.3.3 (metric 3) from 3.3.3.3 (Remote Id:3.3.3.3)
      Origin IGP, metric 0, localpref 100, aigp metric 102      valid, external, best, source safi: 1
      Not advertised to any peer
      Total AIGP metric 105
      Last update: Wed Jun 22 06:57:49 2022

  AS path:100
   Nexthop:2.2.2.2 (metric 2) from 2.2.2.2 (Remote Id:2.2.2.2)
      Origin IGP, metric 0, localpref 100, aigp metric 104      valid, internal, source safi: 1
      Not advertised to any peer
      Total AIGP metric 106
      Last update: Wed Jun 22 06:57:53 2022

R4#
```

Extended Community Attribute

Follow the base configuration as mentioned above.

R1

Route Map Configuration:

`(config)# route-map map1 permit 10`	Configure route-map map1 with sequence number 10.
`(config-route-map)#set aigp-metric 100`	Set aigp metric as 100.
`(config)# route-map map2 permit 20`	Configure route-map map2 with sequence number 20.
`(config-route-map)#set aigp-metric igp-metric`	Set aigp metric as igp-metric.
`(config-route-map)#commit`	Commit the transaction

© 2026 IP Infusion Inc. Proprietary

BGP Configuration:

(config)# router bgp 100	Enter the BGP configuration mode.
(config-router)# neighbor 3.3.3.3 remote-as 100	Configure neighbor
(config-router)# neighbor 2.2.2.2 remote-as 100	Configure neighbor
(config-router)# neighbor 3.3.3.3 update-source 1.1.1.1	Update loopback address as source
(config-router)# neighbor 2.2.2.2 update-source 1.1.1.1	Update loopback address as source
(config-router)# address-family ipv4 unicast	Enter address family mode.
(config-router)#network 50.50.50.0/24	Advertise the network.
(config-router-af)#neighbor 3.3.3.3 activate	Activate the neighbor.
(config-router-af)#neighbor 2.2.2.2 activate	Activate the neighbor.
(config-router-af)#neighbor 3.3.3.3 route-map map1 out	Attach the route-map map1 for the neighbor 3.3.3.3 in out direction
(config-router-af)#neighbor 2.2.2.2 route-map map2 out	Attach the route-map map2 for the neighbor 2.2.2.2 in out direction
(config-router-af)#neighbor 3.3.3.3 aigp disable	Disable the aigp neighbor.
(config-router-af)#neighbor 2.2.2.2 aigp disable	Disable the aigp neighbor.
config-router-af)#neighbor 2.2.2.2 aigp send cost-community 10 poi igp-cost	Enable extended community attribute for the neighbor.
config-router-af)#neighbor 3.3.3.3 aigp send cost-community 10 poi pre-bestpath	Enable extended community attribute for the neighbor.
(config-router-af)#commit	Commit the transaction.

R2

BGP Configuration:

(config)# router bgp 100	Enter the BGP configuration mode.
(config-router)# neighbor 3.3.3.3 remote-as 100	Configure neighbor
(config-router)# neighbor 1.1.1.1 remote-as 100	Configure neighbor
(config-router)# neighbor 4.4.4.4 remote-as 200	Configure neighbor
(config-router)# neighbor 3.3.3.3 update-source 2.2.2.2	Update loopback address as source
(config-router)# neighbor 1.1.1.1 update-source 2.2.2.2	Update loopback address as source
(config-router)# neighbor 4.4.4.4 update-source 2.2.2.2	Update loopback address as source
(config-router)# neighbor 4.4.4.4 ebgp-multihop 2	Enable multihop on eBGP session.
(config-router)# address-family ipv4 unicast	Enter address family mode.
(config-router-af)#neighbor 3.3.3.3 activate	Activate the neighbor.

© 2026 IP Infusion Inc. Proprietary

(config-router-af)#neighbor 1.1.1.1 activate	Activate the neighbor.
(config-router-af)#neighbor 4.4.4.4 activate	Activate the neighbor.
(config-router-af)#commit	Exit address family mode.

R3

BGP Configuration:

(config)# router bgp 100	Enter the BGP configuration mode.
(config-router)#neighbor 1.1.1.1 remote-as 100	Configure neighbor
(config-router)#neighbor 2.2.2.2 remote-as 100	Configure neighbor
(config-router)#neighbor 4.4.4.4 remote-as 200	Configure neighbor
(config-router)#neighbor 1.1.1.1 update-source 3.3.3.3	Update loopback address as source
(config-router)#neighbor 2.2.2.2 update-source 3.3.3.3	Update loopback address as source
(config-router)#neighbor 4.4.4.4 update-source 3.3.3.3	Update loopback address as source
(config-router)# neighbor 4.4.4.4 ebgp-multihop 2	Enable multihop on eBGP session.
(config-router)#address-family ipv4	Enter address family mode.
(config-router-af)#neighbor 1.1.1.1 activate	Activate the neighbor.
(config-router-af)#neighbor 2.2.2.2 activate	Activate the neighbor.
(config-router-af)#neighbor 4.4.4.4 activate	Activate the neighbor.
(config-router-af)#commit	Exit address family mode.

R4

BGP Configuration:

(config)# router bgp 200	Enter the BGP configuration mode.
(config-router)#neighbor 3.3.3.3 remote-as 100	Configure neighbor
(config-router)#neighbor 2.2.2.2 remote-as 100	Configure neighbor
(config-router)# neighbor 3.3.3.3 update-source 4.4.4.4	Update loopback address as source
(config-router)# neighbor 2.2.2.2 update-source 4.4.4.4	Update loopback address as source
(config-router)# neighbor 3.3.3.3 ebgp-multihop 2	Enable multihop on eBGP session.
(config-router)# neighbor 2.2.2.2 ebgp-multihop 2	Enable multihop on eBGP session.
(config-router)#address-family ipv4	Enter address family mode.
(config-router-af)#neighbor 3.3.3.3 activate	Activate the neighbor.
(config-router-af)#neighbor 2.2.2.2 activate	Activate the neighbor.
(config-router-af)#commit	Commit the transaction.

© 2026 IP Infusion Inc. Proprietary

Validation

```
R4#show ip bgp summary
BGP router identifier 4.4.4.4, local AS number 200
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
2.2.2.2                  4   100  164        157       4      0      0  00:02:43               1
3.3.3.3                  4   100  177        194       4      0      0  00:00:28               1

Total number of neighbors 2

Total number of Established sessions 2
R4#show ip bgp
BGP table version is 4, local router ID is 4.4.4.4
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
*>   50.50.50.0/24    2.2.2.2              0        100       0   100 i
*                     3.3.3.3              0        100       0   100 i

Total number of prefixes 1

R4#show ip bgp 50.50.50.0
BGP routing table entry for 50.50.50.0/24
Paths: (2 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  3.3.3.3
  AS path:100
   Nexthop:2.2.2.2 (metric 2) from 2.2.2.2 (Remote Id:2.2.2.2)
      Origin IGP, metric 0, localpref 100      valid, external, best, source safi: 1
      Extended Community:
                         cost(transitive)igp_metric:10:0
      Not advertised to any peer
      Last update: Wed Jun 22 17:57:32 2022

  AS path:100
   Nexthop:3.3.3.3 (metric 3) from 3.3.3.3 (Remote Id:3.3.3.3)
      Origin IGP, metric 0, localpref 100      valid, external, source safi: 1
      Extended Community:
                         cost(transitive)pre-bestpath:10:200
      Not advertised to any peer
       Last update: Wed Jun 22 17:59:46 2022
```

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.

Note: For seamless recovery and optimal functionality, configure VPNv4 GR only for a peer along with VPN AF in BGP. No other BGP capabilities must be configured.

Topology

Figure 64. Device topology for BGP in VR/VRF

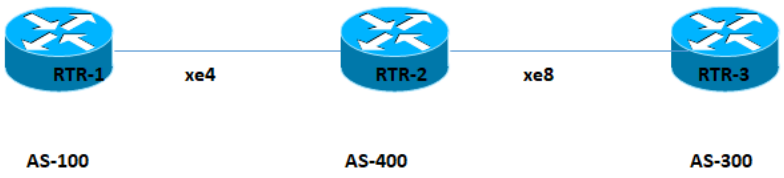

RTR1

`#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 1.1.1.1`	Configure bgp router-id same as loopback ip address.

© 2026 IP Infusion Inc. Proprietary

(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)# exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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)#commit	Commit the candidate configuration to the running configuration.
(config)#router bgp 400	Enter router BGP mode.
(config-router)# bgp router-id 2.2.2.2	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)# exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

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 3.3.3.3	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)# exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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   *>   1.1.1.1/32 is directly connected, lo, 00:10:23
B   *>   2.2.2.2/32 [20/0] via 10.10.10.2, xe4, 00:03:56
B   *>   3.3.3.3/32 [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   *>   1.1.1.1/32 [20/0] via 10.10.10.1, xe4, 00:03:52
C   *>   2.2.2.2/32 is directly connected, lo, 00:07:36
B   *>   3.3.3.3/32 [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.)
```

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   *> 1.1.1.1/32 [20/0] via 20.20.20.1, xe8, 00:01:15
B   *> 2.2.2.2/32 [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.)
```

Validation After BGP Graceful Restart

RTR2

```
#write
Building configuration... [OK]

#restart bgp graceful
%Warning : BGP process will stop and needs to restart manually,
You may lose bgp configuration,if not saved
Proceed for graceful restart? (y/n):y
%% Managed module is down or crashed


#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 BGP   ACTIVE   57      2001/06/07
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

   Network   Next Hop   Metric   LocPrf   Weight   Path
*>   1.1.1.1/32
0.0.0.0   0   100   32768   ?
S>   2.2.2.2/32
10.10.10.2   0   100   0   400   ?
S>   3.3.3.3/32
10.10.10.2   0   100   0   400   300   ?
*>   10.10.10.0/24   0.0.0.0   0   100   32768   ?
S      10.10.10.2   0   100   0   400   ?
S>   20.20.20.0/24   10.10.10.2   0   100   0   400   ?

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

   Network   Next Hop   Metric   LocPrf   Weight   Path
S>   1.1.1.1/32
20.20.20.1   0   100   0   400 100   ?
S>   2.2.2.2/32
20.20.20.1   0   100   0   400 ?
*>   3.3.3.3/32
0.0.0.0   0   100   32768   ?
S>   10.10.10.0/24   20.20.20.1   0   100   0   400 ?
*>   20.20.20.0/24   0.0.0.0   0   100   32768   ?
S      20.20.20.1   0   100   0   400 ?

Total number of prefixes 5
```

BGP Automated Dynamic Route Policy Update

Overview

The Border Gateway Protocol (BGP) is a peer routing protocol, and route maps within BGP play an essential role in filtering and modifying routing information to ensure that only the most optimal routes are advertised. BGP route maps use a combination of prefix lists and sequence numbers, which provide precise control over which IP addresses are allowed or denied. These route maps can be applied in both inbound (Adj-RIB-In) and outbound (Adj-RIB-Out) directions for routing updates to BGP peer neighbors.

In OcNOS 6.5.x or earlier releases, route maps are not automatically applied to a BGP peer neighbor upon updates. Manually executing commands, such as clear ip bgp A.B.C.D (page 1309) with `soft in/out`, is required for each address family to apply route map changes. This manual intervention is necessary to reflect the changes in the BGP Routing Information Base (RIB).

To address this limitation, a new CLI bgp auto-policy-soft-reset enable (page 619), has been introduced. This functionality automatically applies changes to the BGP RIB when route maps or associated lists (such as prefix-list, distribute-list, as-list, access-list, community-list, extended community list, redistribution settings, and more) are applied or updated to a BGP peer neighbor. This dynamic route policy update eliminates the need for the manual execution of the `clear ip bgp <> soft in/out` command for various address families. This greatly simplifies network operations.

The dynamic seamless BGP route policy update introduces the following two new CLIs:

- bgp auto-policy-soft-reset enable (page 619)
- debug bgp soft-clear (page 620)

Notes:

- The bgp auto-policy-soft-reset enable (page 619) functionality is applicable only when apply or remove or update the following:
  - route map
  - prefix-list
  - distribute-list
  - as-list
  - access-list
  - community-list
- This functionality is not supported:
  - when the peer does not support route refresh messages.
  - when the peer is configured with `bgp soft-reconfig-backup`.
  - when the router config disables auto soft reset and the administrator prefers manual resets instead.

Benefits

This feature benefits network administrators by reducing the manual effort involved in BGP route filtering and advertisement updates, improving operational efficiency, and ensuring that routing policies are applied immediately

© 2026 IP Infusion Inc. Proprietary

without requiring a BGP session reset in real-time propagation of routing information.

Prerequisites

Make sure the provider edge routers and BGP route reflectors have the following active configurations.

- Stable BGP session should exist between Provider edge nodes.
- Support for IPv4, IPv6, VPNv4, BGP Labeled-Unicast, VPNv6 address families on all nodes.
- Established BGP neighbor peers or peer groups are up.
- Ensure baseline BGP configurations are proper.
- Check advertised and received routes for accuracy.

Configuring BGP auto-policy-soft-reset

This procedure outlines the steps to configure the bgp auto-policy-soft-reset enable (page 619) functionality, which ensures routing policy changes are applied without disrupting active BGP sessions.

Topology

The configuration uses the below topology, which involves two Provider Edge routers PE1, PE2 and an iBGP router act as a Route Reflector (RR) between PE1 and PE2.

Figure 65. Topology of BGP Auto Policy Soft Reset

Note: Before configuration meet all Prerequisites (page 609)

Procedure

This procedure uses the following interfaces and IP addresses:

- PE1 - `xe3` interface IP address 10.11.1.1/29, loopback IP - 1.1.1.1/32
- PE2 - `ge14` IP address 10.12.3.2/29, loopback IP - 3.3.3.3/32
- RR - `ce46/1` IP address 10.11.1.2/29, loopback IP - 11.11.11.11/32

1. Verify the base BGP configuration on PE1, PE2, and RR before enabling the bgp soft-clear CLI. Refer to the sample out shown in PE1 - Before enabling the bgp auto-policy-soft-reset CLI (page 613), RR - Before enabling the bgp auto-policy-soft-reset CLI (page 614), and PE2 - Before enabling the bgp auto-policy-soft-reset CLI (page 615).
2. Verify the interface, label-switching, IGP, LDP, existing route maps filter criteria such as prefix-list on PE1 and PE2. Here, is the sample configuration.

   PE1

   ```
   debug bgp soft-clear
   .
   ```

```
.
ip vrf management
!
ip vrf vrf221
 rd 221:1
 route-target both 221:1
!
!
router ldp
 session-protection
 targeted-peer ipv4 3.3.3.3
  exit-targeted-peer-mode
 transport-address ipv4 1.1.1.1
!

interface lo
 ip address 127.0.0.1/8
 ip address 1.1.1.1/32 secondary
 ipv6 address ::1/128
!
interface xe3
 description connected-to-RR
 load-interval 30
  ip address 10.11.1.1/29
 ipv6 address 3001::2/64
 mtu 9216
 label-switching
 ipv6 router ospf area 0.0.0.0 instance-id 0
 enable-ldp ipv4
!
router ospf 100
 ospf router-id 1.1.1.1
 network 1.1.1.1/32 area 0.0.0.0
 network 10.2.1.0/29 area 0.0.0.0
 network 10.11.1.0/29 area 0.0.0.0
!
router ipv6 ospf
 router-id 1.1.1.1
!
```

PE2

```
debug bgp soft-clear
 ip vrf vrf221
 rd 221:1
 route-target both 221:1
!
router ldp
 auto-targeted-session
 session-protection
 transport-address ipv4 3.3.3.3
!
interface lo
 ip address 127.0.0.1/8
 ip address 3.3.3.3/32 secondary
 ipv6 address ::1/128
!
interface ge14
 description connected-to-RR
 load-interval 30
 ip address 10.12.3.2/29
 ipv6 address 6001::3/64
 mtu 9216
 label-switching
 ipv6 router ospf area 0.0.0.0 instance-id 0
 enable-ldp ipv4
!
router ospf 100
```

```
 network 3.3.3.3/32 area 0.0.0.0
 network 10.12.3.0/29 area 0.0.0.0
 network 10.12.4.0/29 area 0.0.0.0
!
router ipv6 ospf
 router-id 3.3.3.3
!
```

RR

```
debug bgp soft-clear
!
ip vrf management
!
router ldp
!
interface ce46/1
 description connected-to-PE1
 load-interval 30
 ip address 10.11.1.2/29
 ipv6 address 3001::3/64
 mtu 9216
 label-switching
 ipv6 router ospf area 0.0.0.0 instance-id 0
 enable-ldp ipv4
!
interface lo
 ip address 127.0.0.1/8
 ip address 11.11.11.11/32 secondary
 ipv6 address ::1/128
!
interface xe3
 description connected-to-PE2
 speed 1g
 load-interval 30
 ip address 10.12.3.3/29
 ipv6 address 6001::2/64
 mtu 9216
 label-switching
 ipv6 router ospf area 0.0.0.0 instance-id 0
 enable-ldp ipv4
!
 exit
!
router ospf 100
 ospf router-id 11.11.11.11
 network 10.11.1.0/29 area 0.0.0.0
 network 10.11.12.0/29 area 0.0.0.0
 network 10.12.3.0/29 area 0.0.0.0
 network 11.11.11.11/32 area 0.0.0.0
!
router ipv6 ospf
 router-id 11.11.11.11
!
```

3. Enable `bgp auto-policy-soft-reset` on PE1, PE2, and RR. Here, is the sample configuration.

   PE1

```
!
router bgp 100
 bgp router-id 1.1.1.1
  bgp auto-policy-soft-reset enable
 !
```

   PE2

```
!
router bgp 100
 bgp router-id 3.3.3.3
  bgp auto-policy-soft-reset enable
 !
```

RR

```
!
router bgp 100
 bgp router-id 11.11.11.11
  bgp auto-policy-soft-reset enable
 no bgp inbound-route-filter
 neighbor peer_group1 peer-group
 neighbor peer_group1 remote-as 100
 neighbor peer_group1 update-source lo
 neighbor peer_group1 advertisement-interval 0
  neighbor 1.1.1.1 peer-group peer_group1
 neighbor 3.3.3.3 peer-group peer_group1
 !
```

4. Add a route map policy on PE1. For example, configure route map to allow or deny additional prefixes and change the attributes values.

```
!
ip prefix-list prefix1
 seq 5 permit 31.0.1.0/24
!
route-map rm_pe1 permit 1
 match ip address prefix-list prefix1
 set metric 200
 set weight 200
 set local-preference 200
 set community 100:101
!
```

5. Add a route map policy on PE2. For example, configure route map to allow or deny additional prefixes and change the attributes values.

```
!
ip community-list standard Community1 permit 100:101
!
route-map rm_pe2 permit 1
 match community Community1
 set community 100:102
!
```

6. Attach the updated route map filter to BGP peers on PE1 and PE2. For example, on PE1, attach the route map to filter specific routes on outgoing routes using "`neighbor 11.11.11.11 route-map rm-pe1 out`" command towards PE2 under address family vpnv4 unicast. On, PE2, attach it to filter specific routes on incoming routes using "`neighbor 11.11.11.11 route-map rm-pe2 in`" command under address family vpnv4 unicast.

PE1

```
!
router bgp 100
 bgp router-id 1.1.1.1
!
address-family vpnv4 unicast
 neighbor 11.11.11.11 activate
 neighbor 11.11.11.11 route-map rm_pe1 out
 exit-address-family
!
```

PE2

```
!
router bgp 100
bgp router-id 3.3.3.3
!
address-family vpnv4 unicast
 neighbor 11.11.11.11 activate
 neighbor 11.11.11.11 route-map rm_pe2 in
 exit-address-family
!
```

7. Verify the advertised and received routes to ensure the policy changes have taken effect on PE1, PE2 and RR. Refer to the sample out shown in PE2 - After enabling the bgp auto-policy-soft-reset (page 617).

Validation

PE1 - Before enabling the bgp auto-policy-soft-reset CLI

The following show output displays the active BGP configuration PE1

```
PE1#show running-config bgp
!
router bgp 100
 bgp router-id 1.1.1.1
 neighbor 11.11.11.11 remote-as 100
 neighbor 11.11.11.11 update-source lo
 neighbor 11.11.11.11 advertisement-interval 0
 !
 address-family ipv4 unicast
 network 1.1.1.1/32
 exit-address-family
 !
 address-family ipv4 labeled-unicast
 neighbor 11.11.11.11 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 11.11.11.11 activate
 exit-address-family
 !
 address-family vpnv6 unicast
 neighbor 11.11.11.11 activate
 exit-address-family
 !
 address-family ipv6 unicast
 neighbor 11.11.11.11 activate
 exit-address-family
 !
 address-family ipv4 vrf vrf221
 redistribute connected
 redistribute ospf
 neighbor 101.1.1.10 remote-as 200
 neighbor 101.1.1.10 activate
 exit-address-family
 !
 address-family ipv6 vrf vrf221
 redistribute connected
 exit-address-family
 !
 exit
!
PE1#
PE1#show ip bgp labeled-unicast summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 1
2 BGP AS-PATH entries
0 BGP community entries
```

```
Neighbor           V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
11.11.11.11        4   100         49         36       1      0      0  00:11:55             0

Total number of neighbors 1

Total number of Established sessions 1
PE1#
PE1#show ip bgp vpnv4 vrf vrf221
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
Route Distinguisher: 221:1 (Default for VRF vrf221)
*> l 31.0.0.0/24      101.1.1.10            0        100        0   200 i
*> l 31.0.1.0/24      101.1.1.10            0        100        0   200 i
*> l 31.0.2.0/24      101.1.1.10            0        100        0   200 i
*> l 33.0.0.0/24      101.1.1.10            0        100        0   200 i
*> l 33.0.1.0/24      101.1.1.10            0        100        0   200 i
*> l 33.0.2.0/24      101.1.1.10            0        100        0   200 i
*>il 33.0.3.0/24      3.3.3.3               0        100        0   200 i
*>il 33.0.4.0/24      3.3.3.3               0        100        0   200 i
*> l 101.1.1.0/24     0.0.0.0               0        100        32768  ?
*>il 151.1.1.0/24     3.3.3.3               0        100        0    ?
 Announced routes count = 7
 Accepted routes count = 3
PE1#
PE1#show ip bgp vpnv4 vrf vrf221 31.0.1.0
Route Distinguisher: 221:1 (Default for VRF vrf221) Routing Entry for prefix: 31.0.1.0/24
  Advertised to non peer-group peers:
  11.11.11.11 11.11.11.11
  AS path:200
  Path Selection reason: Nothing left to compare
   Nexthop:101.1.1.10 (IGP metric 0) from 101.1.1.10 (Remote Id:192.0.4.53) Peer nexthop: 101.1.1.10
      Origin IGP, metric 0, localpref 100, Out-label 0, In-label 24384, refcnt: 1
      valid, external, best, source-safi: 1
      Extended Community: RT:221:1
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Last update: Fri Dec  6 18:11:29 2024, 00:12:28 ago
```

RR - Before enabling the bgp auto-policy-soft-reset CLI

The following show output displays the active BGP configuration on RR b

```
RR#show running-config bgp
!
router bgp 100
 bgp router-id 11.11.11.11
 bgp auto-policy-soft-reset enable
 no bgp inbound-route-filter
 neighbor peer_group1 peer-group
 neighbor peer_group1 remote-as 100
 neighbor peer_group1 update-source lo
 neighbor peer_group1 advertisement-interval 0
 neighbor 1.1.1.1 peer-group peer_group1
 neighbor 3.3.3.3 peer-group peer_group1
 !
 address-family ipv4 unicast
 network 11.11.11.11/32
 exit-address-family
 !
 address-family ipv4 labeled-unicast
 neighbor peer_group1 activate
 neighbor peer_group1 route-reflector-client
```

```
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor peer_group1 activate
 neighbor peer_group1 route-reflector-client
 exit-address-family
 !
 address-family vpnv6 unicast
 neighbor peer_group1 activate
 neighbor peer_group1 route-reflector-client
 exit-address-family
 !
 address-family ipv6 unicast
 neighbor peer_group1 activate
 neighbor peer_group1 route-reflector-client
 exit-address-family
 !
 exit
!
RR#
RR#show ip bgp labeled-unicast summary
BGP router identifier 11.11.11.11, local AS number 100
BGP table version is 1
2 BGP AS-PATH entries
0 BGP community entries

Neighbor           V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
1.1.1.1            4   100        124        142       1      0      0  00:09:33             0
3.3.3.3            4   100        140        131       1      0      0  00:01:14             0

Total number of neighbors 2

Total number of Established sessions 2
RR#
RR#show ip bgp vpnv4 all summary
BGP router identifier 11.11.11.11, local AS number 100
BGP table version is 14
2 BGP AS-PATH entries
0 BGP community entries

Neighbor           V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
1.1.1.1            4   100        124        142      14      0      0  00:09:47             7
3.3.3.3            4   100        140        131      14      0      0  00:01:28             3

Total number of neighbors 2

Total number of Established sessions 2
RR#
```

PE2 - Before enabling the bgp auto-policy-soft-reset CLI

The following show output displays the active BGP configuration on PE2.

```
PE2#show running bgp
!
router bgp 100
 bgp router-id 3.3.3.3
 bgp auto-policy-soft-reset enable
 neighbor 11.11.11.11 remote-as 100
 neighbor 11.11.11.11 update-source lo
 neighbor 11.11.11.11 advertisement-interval 0
 !
 address-family ipv4 unicast
 network 3.3.3.3/32
 exit-address-family
 !
```

```
 address-family ipv4 labeled-unicast
 neighbor 11.11.11.11 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 11.11.11.11 activate
 neighbor 11.11.11.11 route-map rm_pe2 in
 exit-address-family
 !
 address-family vpnv6 unicast
 neighbor 11.11.11.11 activate
 exit-address-family
 !
 address-family ipv6 unicast
 exit-address-family
 !
 address-family ipv4 vrf vrf221
 redistribute connected
 redistribute ospf
 neighbor 151.1.1.10 remote-as 200
 neighbor 151.1.1.10 activate
 exit-address-family
 !
 address-family ipv6 vrf vrf221
 redistribute connected
 exit-address-family
 !
 exit
!

PE2#show ip bgp labeled-unicast summary
BGP router identifier 3.3.3.3, local AS number 100
BGP table version is 1
2 BGP AS-PATH entries
0 BGP community entries

Neighbor           V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
11.11.11.11        4   100         45         55       1      0      0  00:05:17             0

Total number of neighbors 1

Total number of Established sessions 1

PE2#show ip bgp vpnv4 vrf vrf221
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
Route Distinguisher: 221:1 (Default for VRF vrf221)
*>il 31.0.0.0/24      1.1.1.1               0        100        0   200 i
*>il 31.0.1.0/24      1.1.1.1               0        100        0   200 i
*>il 31.0.2.0/24      1.1.1.1               0        100        0   200 i
*>il 33.0.0.0/24      1.1.1.1               0        100        0   200 i
*>il 33.0.1.0/24      1.1.1.1               0        100        0   200 i
*>il 33.0.2.0/24      1.1.1.1               0        100        0   200 i
*> l 33.0.3.0/24      151.1.1.10            0        100        0   200 i
*> l 33.0.4.0/24      151.1.1.10            0        100        0   200 i
*>il 101.1.1.0/24     1.1.1.1               0        100        0    ?
*> l 151.1.1.0/24     0.0.0.0               0        100        32768  ?
 Announced routes count = 3
 Accepted routes count = 7
PE2#
PE2#show ip bgp vpnv4 vrf vrf221 31.0.1.0
Route Distinguisher: 221:1 (Default for VRF vrf221) Routing Entry for prefix: 31.0.1.0/24
  Not advertised to any peer
  AS path:200
```

```
  Path Selection reason: Nothing left to compare
 Nexthop:1.1.1.1 (IGP metric 3) from 11.11.11.11 (Originator Id:1.1.1.1) (Remote Id:11.11.11.11) Peer
nexthop: 11.11.11.11

 Origin IGP, metric 0, localpref 100, Out-label 24384      valid, internal, best, source-safi: 128

      Duplicated: (source VRF-ID: 0, source VRF: DEFAULT, VRF-External, imported)
      Extended Community: RT:221:1
      Originator: 1.1.1.1, Cluster list: 11.11.11.11
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Last update: Fri Dec  6 18:19:48 2024, 00:05:25 ago
```

PE1 - After enabling the bgp auto-policy-soft-reset

Following show output displays the modified route policy.

```
PE1#show running-config route-map
!
route-map rm_pe1 permit 1
 match ip address prefix-list prefix1
 set metric 200
 set weight 200
 set local-preference 200
 set community 100:101
!
PE1#
```

PE2 - After enabling the bgp auto-policy-soft-reset

The following show output displays the active BGP configuration on PE2.

```
PE2#show ip bgp vpnv4 vrf vrf221 31.0.1.0
Route Distinguisher: 221:1 (Default for VRF vrf221) Routing Entry for prefix: 31.0.1.0/24
  Not advertised to any peer
  AS path:200
  Path Selection reason: Nothing left to compare
 Nexthop:1.1.1.1 (IGP metric 3) from 11.11.11.11 (Originator Id:1.1.1.1) (Remote Id:11.11.11.11) Peer
nexthop: 11.11.11.11

      Origin IGP, metric 200, localpref 200, Out-label 24384      valid, internal, best, source-safi:
128
  Duplicated: (source VRF-ID: 0, source VRF: DEFAULT, VRF-External, imported)

      Community: 100:102
 Extended Community: RT:221:1

      Originator: 1.1.1.1, Cluster list: 11.11.11.11
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Last update: Mon Dec  2 23:18:45 2024, 00:17:31 ago

PE2#
```

Implementation Examples

Following are some of the use cases where manual clearance clear ip bgp A.B.C.D (page 1309) of route updates are necessary before applying them to a neighbor peer or peer group.

Scenario 1:

When configuring bgp additional-paths (page 1595) with options send|receive|send-receive|select (all | best<2-3>))or when enabling neighbor advertise additional-

paths (page 1598), route change occurs that needs to be updated in the neighbor peer or peer group through BGP RIB-out.

Scenario 2:

Configuring neighbor unsuppress-map (page 1404) leads to route changes that needs to be announced in BGP RIB-out.

Scenario 3:

Enabling neighbor allowas-in (page 1340) causes change in routes that needs to accepted via BGP RIB-in.

Scenario 4:

When multiple BGP address families facilitate route management between two customer edge devices (CE1 and CE2) via Route Reflector, enabling bgp auto-policy-soft-reset enable (page 619) dynamically updates newly announced route map prefixes.

CLI Commands

Following are the CLIs introduced to clear the route-map updates automatically for each BGP address-family.

bgp auto-policy-soft-reset enable

Use this command ONLY to enable BGP soft clear for address families IPv4, IPv6, VPNv4, VPNv6, labeled-unicast regarding route map, table-map, prefix-list, distribute-list, as-list, access-list or community-list updates to apply the modified routes and attributes automatically on the neighbor peer or peer group.

Use the `no` form of this CLI to disable the soft reset.

Note: EVPN address family is not supported.

Command Syntax

```
bgp auto-policy-soft-reset enable
no bgp auto-policy-soft-reset enable
```

Parameters

None

Default

Disabled

Command Mode

BGP Router Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

Entering the CLI enables the soft clear the route map.

```
OcNOS(config)#router bgp
OcNOS(config-router)#bgp auto-policy-soft-reset enable
OcNOS(config-router)#address-family vpnv4 unicast
OcNOS(config-router-af)#neighbor 50.50.50.50 prefix-list prefix-inside in
OcNOS(config-router-af)#commit
OcNOS(config-router-af)#exit
OcNOS(config-router)#commit
```

© 2026 IP Infusion Inc. Proprietary

debug bgp soft-clear

Use this command to log the BGP auto soft clear information with timestamp for debug purpose.

Command Syntax

```
debug bgp soft-clear
```

Parameters

None

Default

Disabled

Command Mode

Configure Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

Entering the CLI enables the BGP soft clear log file.

```
OcNOS(config)#debug bgp soft-clear
```

Glossary

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

Key Terms/Acronym	Description
AS	Autonomous System
BGP	Border Gateway Protocol router
PE	Provider Edge router
RIB	Route Information Base
RR	Route-Reflector router
Route-map	Maps used for filtering RIB in/out on various use cases
Prefix list	List created on router with Network prefixes ge, le options too to indicate network address length
Access List	List created on a router with IP addresses mentioned (IPv4, IPv6) without any network address length flexibility like prefix lists.
AS List	Autonomous System path list where BGP can filter routes based on AS path it checks in the control packet
Community List	Community is a mini group inside AS, and BGP can filter these communities based on AS:NN where the first 16 bits are AS, next 16 bits a unique number with AS.

© 2026 IP Infusion Inc. Proprietary

BGP4+

This section contains basic BGP4+ configuration examples.

For details about the commands used in these examples, see the Border Gateway Protocol Command Reference.

Enable iBGP Peering Using a Global Address

This example shows the minimum configuration required to enable BGP on an interface. R1 and R2 are two routers belonging to the same Autonomous System (AS), AS200, connecting to network **`3ffe:10::/48`**. First, specify the IPv6 global address, then define the routing process and AS number to which the routers belong. Configure a fixed Router ID, then, define BGP neighbors to start exchanging routing updates.

Topology

Figure 66. iBGP Peering

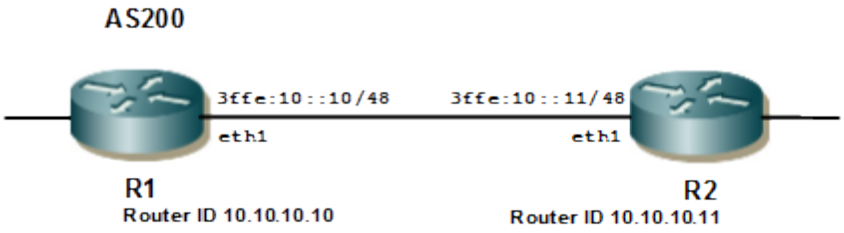

Configuration

R1

`#configure terminal`	Enter Configure mode.
`(config)#interface eth1`	Specify the interface `(eth1)` to configure.
`(config-if)#ipv6 address 3ffe:10::10/48`	Specify the IPv6 global address.
`(config-if)#exit`	Enter Configure mode.
`(config)#router bgp 200`	Define the routing process. The number `200` specifies the AS number of `R2`.
`(config-router)#bgp router-id 10.10.10.11`	Configure a fixed Router ID (`10.10.10.11`) for the BGP4+ routing process.
`(config-router)#neighbor 3ffe:10::10 remote-as 200`	Define the BGP neighbor (R1), and establish a TCP session by specifying the global IPv6 address (`3ffe:10::10`) and the AS number `(200)` of neighbor `R1`.
`(config-router)#address-family ipv6`	Enter Address Family mode for configuring routing sessions that use IPv6 address prefixes.
`(config-router-af)#neighbor 3ffe:10::10 activate`	Activate the neighbor `R1(3ffe:10::10)`, and

© 2026 IP Infusion Inc. Proprietary

	enable exchange of IPv6 address prefix types with this neighbor.
(config-router-af)#exit-address-family	Exit address family
(config-router)#commit	Commit the candidate configuration to the running configuration

R2

#configure terminal	Enter Configure mode.
(config)#interface eth1	Specify the interface (eth1) to configure.
(config-if)#ipv6 address 3ffe:10::11/48	Specify the IPv6 global address.
(config-if)#exit	Enter Configure mode.
(config)#router bgp 200	Define the routing process. The number 200 specifies the AS number of R2.
(config-router)#bgp router-id 10.10.10.11	Configure a fixed Router ID (10.10.10.11) for the BGP4+ routing process.
(config-router)#neighbor 3ffe:10::10 remote-as 200	Define the BGP neighbor (R1), and establish a TCP session by specifying the global IPv6 address (3ffe:10::10) and the AS number (200) of neighbor R1.
(config-router)#address-family ipv6	Enter Address Family mode for configuring routing sessions that use IPv6 address prefixes.
(config-router-af)#neighbor 3ffe:10::10 activate	Activate the neighbor R1(3ffe:10::10), and enable exchange of IPv6 address prefix types with this neighbor.
(config-router-af)#exit-address-family	Exit address family
(config-router)#commit	Commit the candidate configuration to the running configuration

Validation

show bgp ipv6 summary, show bgp ipv6 neighbors.

R1

```
#show bgp ipv6 summary
BGP router identifier 10.10.10.10, 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
3ffe:10::11              4   200   10         10       1      0      0  00:03:36              0

Total number of neighbors 1

Total number of Established sessions 1
```

```
#show bgp ipv6 neighbors
BGP neighbor is 3ffe:10::11, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 10.10.10.10, remote router ID 10.10.10.11
  BGP state = Established, up for 00:05:39
  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 IPv6 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: IPv6 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: 3ffe:10::10, Local port: 44754
Foreign host: 3ffe:10::11, Foreign port: 179
Nexthop: 10.10.10.10
Nexthop global: 3ffe:10::10
Nexthop local: fe80::5054:ff:fe11:d4f5
BGP connection: shared network
```

R2

```
R2#show bgp ipv6 summary
BGP router identifier 10.10.10.11, 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
3ffe:10::10              4   200   31         32       1      0      0  00:12:53               0

Total number of neighbors 1

Total number of Established sessions 1

R2#show bgp ipv6 neighbors
BGP neighbor is 3ffe:10::10, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 10.10.10.11, remote router ID 10.10.10.10
  BGP state = Established, up for 00:15:29
  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 IPv6 Unicast: advertised and received
  Received 37 messages, 0 notifications, 0 in queue
  Sent 39 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 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: 3ffe:10::11, Local port: 179
Foreign host: 3ffe:10::10, Foreign port: 44754
Nexthop: 10.10.10.11
```

```
Nexthop global: 3ffe:10::11
Nexthop local: fe80::5054:ff:fe82:8dcc
BGP connection: shared network
```

Enable iBGP Peering Using Link-local Address

This example shows the minimum configuration required to enable iBGP on an interface. R1 and R2 are two routers belonging to the same AS, AS200, connecting to network fe80::/10. First, define the routing process and AS number to which the routers belong. Configure a fixed Router ID for the BGP4+ routing process, then, define BGP neighbors to start exchanging routing updates.

Topology

Figure 67. iBGP Peering Link-Local Address

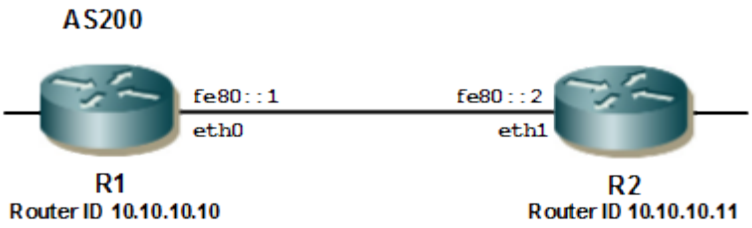

Configuration

R1

`#configure terminal`
 	
`(config)#router bgp 200`
	Enter Configure mode.
	
	Define the routing process. The number `200` specifies the AS number of `R1`.

`(config-router)#bgp router-id 10.10.10.10`	Configure a fixed Router ID (`10.10.10.10`) for the BGP4+ routing process.
`(config-router)#neighbor fe80::2 remote-as 200`	Define BGP neighbor (R2), and establish a TCP session by specifying the link-local address (`fe80::2`) and the AS number `(200)` of neighbor `R2`.
`(config-router)#neighbor fe80::2 interface eth0`	To specify a link-local neighbor, configure the interface name of the neighbor `fe80::2`.
`(config-router)#address-family ipv6`	Enter Address Family mode for configuring routing sessions that use IPv6 address prefixes.
`(config-router-af)#neighbor fe80::2 activate`	Activate the neighbor `R2(fe80::2)`, and enable exchange of IPv6 address prefix types with this neighbor.
`(config-router-af)#exit-address-family`	Exit address family
`(config-router)#commit`	Commit the candidate configuration to the running configuration

© 2026 IP Infusion Inc. Proprietary

R2

`(config)#router bgp 200`	Define the routing process. The number `200` specifies the AS number of `R2`.
`(config-router)#bgp router-id 10.10.10.11`	Configure a fixed Router ID (`10.10.10.11`) for the BGP4+ routing process.
`(config-router)#neighbor fe80::1 remote-as 200`	Define the BGP neighbor (R1), and establish a TCP session by specifying the link-local address `R1` (`fe80::1`) and the AS number `(200)` of neighbor `R1`.
`(config-router)#neighbor fe80::1 interface eth1`	To specify a link-local neighbor, configure the interface name of the neighbor `fe80::1`.
`(config-router)#address-family ipv6`	Enter Address Family mode for configuring routing sessions that use IPv6 address prefixes.
`(config-router-af)#neighbor fe80::1 activate`	Activate the neighbor `R1(fe80:1)`, and enable exchange of IPv6 address prefix types with this neighbor.
`(config-router-af)#exit-address-family`	Exit address family
`(config-router)#commit`	Commit the candidate configuration to the running configuration

Validation

show bgp ipv6 summary, show bgp ipv6, show bgp ipv6 neighbors

R1

```
R1#show bgp ipv6 summary
BGP router identifier 10.10.10.10, 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
fe80::2  4   200   68         67       1      0      0  00:27:49               0

Total number of neighbors 1

Total number of Established sessions 1
R1#show bgp ipv6 neighbors
BGP neighbor is fe80::2, remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 10.10.10.10, remote router ID 10.10.10.11
  BGP state = Established, up for 00:28:14
  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 IPv6 Unicast: advertised and received
  Received 68 messages, 0 notifications, 0 in queue
  Sent 68 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 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::5054:ff:fe11:d4f5, Local port: 179
Foreign host: fe80::2, Foreign port: 44058
Nexthop: 10.10.10.10
Nexthop global: 3ffe:10::10
Nexthop local: fe80::5054:ff:fe11:d4f5
BGP connection: shared network
```

R2

```
R2#show bgp ipv6 summary
BGP router identifier 10.10.10.11, 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
fe80::1   4   200   54         55       1      0      0  00:22:19               0

Total number of neighbors 1

Total number of Established sessions 1

R2#show bgp ipv6 neighbors
BGP neighbor is fe80::1 , remote AS 200, local AS 200, internal link
  BGP version 4, local router ID 10.10.10.11, remote router ID 10.10.10.10
  BGP state = Established, up for 00:22:59
  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 IPv6 Unicast: advertised and received
  Received 55 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: IPv6 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::5054:ff:fe82:8dcc, Local port: 44058
Foreign host: fe80::1 , Foreign port: 179
Nexthop: 10.10.10.11
Nexthop global: 3ffe:10::11
Nexthop local: fe80::5054:ff:fe82:8dcc
BGP connection: shared network
```

Enable eBGP Peering Between Different Autonomous Systems

This example shows the minimum configuration required to enable eBGP on an interface, when the routers belong to different ASs. R1 and R2 are two routers in different ASs, AS200 and AS300 connecting to network 3ffe:10::/64.

Topology

Figure 68. BGP Peering - Different AS

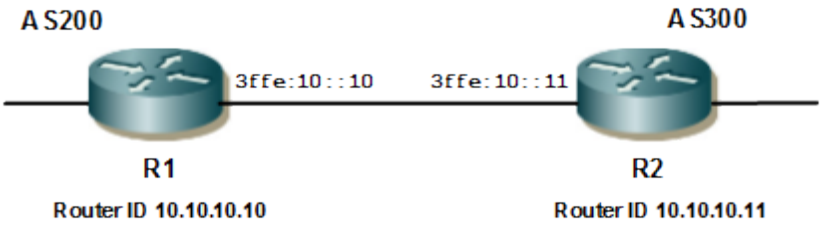

Configuration

R1

`#configure terminal`
 	
`(config)#router bgp 200`
	Enter Configure mode.
	
	Define the routing process. The number `200` specifies the AS number of `R1`.

`(config-router)#bgp router-id 10.10.10.10`	Configure a fixed Router ID (`10.10.10.10`) for the BGP4+ routing process.
`(config-router)#neighbor 3ffe:10::11 remote-as 300`	Define the BGP neighbor (R2), and establish a TCP session by specifying the IPv6 address (`3ffe:10::11`) and the AS number (`300`) of neighbor `R2`.
`(config-router)#address-family ipv6`	Enter Address Family mode for configuring routing sessions that use IPv6 address prefixes.
`(config-router-af)#neighbor 3ffe:10::11 activate`	Activate the neighbor `R2 (3ffe:10::11)`, and enable exchange of IPv6 address prefix types with this neighbor.
`(config-router-af)#exit-address-family`	Exit address family
`(config-router)#commit`	Commit the candidate configuration to the running configuration

R2

`(config)#router bgp 300`	Define the routing process. The number `300` specifies the AS number of `R2`.
`(config-router)#bgp router-id 10.10.10.11`	Configure a fixed Router ID (`10.10.10.11`) for the BGP4+ routing process.
`(config-router)#neighbor 3ffe:10::10 remote-as 200`	Define the BGP neighbor (R1), and establish a TCP session by specifying the IPv6 address (`3ffe:10::10`) and the AS number (`200`) of neighbor `R1`.
`(config-router)#address-family ipv6`	Enter Address Family mode for configuring routing

© 2026 IP Infusion Inc. Proprietary

	sessions that use IPv6 address prefixes.
`(config-router-af)#neighbor 3ffe:10::10 activate`	Activate the neighbor `R1(3ffe:10::10)` and enable exchange of IPv6 address prefix types with this neighbor.
`(config-router-af)#exit-address-family`	Exit address family
`(config-router)#commit`	Commit the candidate configuration to the running configuration

Validation

show bgp ipv6 summary, show bgp ipv6 neighbors.

R1

```
R1#show bgp ipv6 summary
BGP router identifier 10.10.10.10, 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
3ffe:10::11              4   300    3          4       1      0      0  00:00:37               0

Total number of neighbors 1
Total number of Established sessions 1

R1#show bgp ipv6 neighbors
BGP neighbor is 3ffe:10::11, remote AS 300, local AS 200, external link
  BGP version 4, local router ID 10.10.10.10, remote router ID 10.10.10.11
  BGP state = Established, up for 00:01:42
  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 IPv6 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: IPv6 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: 3ffe:10::10, Local port: 179
Foreign host: 3ffe:10::11, Foreign port: 46240
Nexthop: 10.10.10.10
Nexthop global: 3ffe:10::10
Nexthop local: fe80::5054:ff:fe11:d4f5
BGP connection: shared network
```

R2

```
R2#show bgp ipv6 summary
BGP router identifier 10.10.10.11, 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
3ffe:10::10              4   200    2          2       1      0      0  00:00:24               0

Total number of neighbors 1

Total number of Established sessions 1

R2#show bgp ipv6 neighbors
BGP neighbor is 3ffe:10::10, remote AS 200, local AS 300, external link
  BGP version 4, local router ID 10.10.10.11, remote router ID 10.10.10.10
  BGP state = Established, up for 00:01:32
  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 IPv6 Unicast: advertised and received
  Received 5 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: IPv6 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: 3ffe:10::11, Local port: 46240
Foreign host: 3ffe:10::10, Foreign port: 179
Nexthop: 10.10.10.11
Nexthop global: 3ffe:10::11
Nexthop local: fe80::5054:ff:fe82:8dcc
BGP connection: shared network
```

Route-Map

Use route-maps to filter incoming updates from a BGP peer. In this example, a prefix list named myPrefixList on R1 is configured to deny entry of any routes with the IP address **`3ffe:12::/32`**. To test the filter, R2 is configured to generate network prefixes **`3ffe:11::/48`** and **`3ffe:12::/48`**. To verify, use the show bgp ipv6 command on R1; it displays R1 receiving only the **`3ffe:11::/48`** network prefix.

Topology

Figure 69. Route-Map

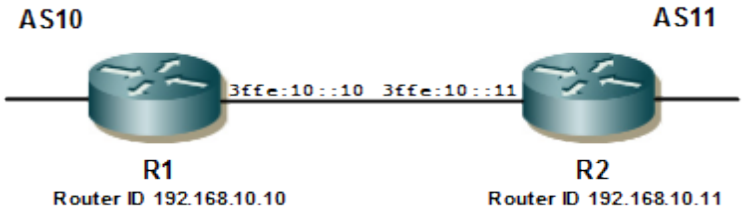

Configuration

R1

`#configure terminal`	Enter Configure mode.
`(config)Interface eth1`	Specify the interface (**`eth1`**) to configure
`(config)Ipv6 address 3ffe:10::10/64`	Specify the IPv6 global address
`(config)Exit`	Enter Configure mode
`(config)#ipv6 prefix-list myPrefixList`	Create an entry in the prefix-list. myPrefixList is the name of the map that is created
`(config-ipv6-prefix-list)#seq 5 deny 3ffe:12::/32 ge 48 le 64`	Here 5 and 10 specify the sequence number or position of this specific route map. deny specifies the packets are to be rejected. permit specifies the packets are to be allowed. 48 and 64 are the minimum and maximum prefix lengths, respectively, to be matched.
`(config-ipv6-prefix-list)#seq 10 permit any`	Create another entry in the myPrefixList map. 10 specifies the sequence number or position of this specific route map. permit any specifies accept all packets of any length.
`(config-ipv6-prefix-list)#exit`	Exit the prefix-list mode
`(config)#route-map myPrefixList permit 1`	Enter Route-map mode.
`(config-route-map)#match ipv6 address prefix-list myPrefixList`	Set the match criteria. In this case, if the route-map name matches myPrefixList, the packets from the first sequence will be denied.
`(config-route-map)#exit`	Exit Route-map mode, and return to Configure mode.
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#router bgp 10`	Define the routing process. The number `10` specifies the AS number of `R1`.
`(config-router)#bgp router-id 192.168.10.10`	Configure a fixed Router ID (`192.168.10.10`) for the BGP4+ routing process.
`(config-router)#neighbor 3ffe:10:11 remote-as 11`	Define the BGP neighbor (R2), and establish a TCP session by specifying the IPv6 address (`3ffe:10::11`) and the AS number (`11`) of neighbor `R2`.
`(config-router-af)#exit`	Exit from router BGP mode.
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#address-family ipv6`	Enter Address Family mode for configuring routing sessions that use IPv6 address prefixes.
`(config-router-af)#neighbor 3ffe:10::11 activate`	Activate the neighbor `R2(3ffe:10::11)`, and

	enable exchange of IPv6 address prefix types with this neighbor.
(config-router-af)#neighbor 3ffe:10::11 route-map myPrefixList in	Apply the route-map myPrefixList to all incoming routes.
(config-router-af)#exit-address-family	Exit Address Family mode, and return to Router mode.
(config-router-af)#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)Interface eth1	Specify the interface (**eth1**) to configure
(config)Ipv6 address 3ffe:10::11/64	Specify the IPv6 global address
(config)Exit	Enter Configure mode
(config)#router bgp 11	Define the routing process. The number 11 specifies the AS number of R2.
(config-router)#bgp router-id 192.168.10.11	Configure a fixed Router ID (192.168.10.11) for the BGP4+ routing process.
(config-router)#neighbor 3ffe:10::10 remote-as 10	Define the BGP neighbor (R1), and establish a TCP session by specifying the IPv6 address (3ffe:10::10) and the AS number (10) of neighbor R1.
(config-router)#address-family ipv6	Enter Address Family mode for configuring routing sessions that use IPv6 address prefixes.
(config-router-af)#network 3ffe:11::/48	Announce the IPv6 network prefix (3ffe:11::/48).
(config-router-af)#network 3ffe:12::/48	Announce the IPv6 network prefix (3ffe:12::/48).
(config-router-af)#neighbor 3ffe:10::10 activate	Activate the neighbor R1(3ffe:10::10), and enable exchange of IPv6 address prefix types with this neighbor.
(config-router-af)#exit-address-family	Exit Address Family mode, and return to Router mode.
(config)#commit	Commit the candidate configuration to the running configuration.

Validation

show bgp ipv6 summary, show bgp ipv6 neighbors, show bgp ipv6, show bgp ipv6 prefix-list.

R1

```
R1#show ipv6 bgp summary
```

```
BGP router identifier 192.168.10.10, local AS number 10
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
3ffe:10::11              4    11  140        306       1      0      0  00:10:44               0

Total number of neighbors 1

Total number of Established sessions 1



R1#show bgp ipv6 neighbors
BGP neighbor is 3ffe:10::11, remote AS 11, local AS 10, external link
  BGP version 4, local router ID 192.168.10.10, remote router ID 192.168.10.11
  BGP state = Established, up for 00:11:02
  Last read 00:00:23, 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 140 messages, 0 notifications, 0 in queue
  Sent 194 messages, 113 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv6 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Inbound path policy configured
  Route map for incoming advertisements is *myPrefixList
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 3ffe:10::10, Local port: 179
Foreign host: 3ffe:10::11, Foreign port: 46796
Nexthop: 192.168.10.10
Nexthop global: 3ffe:10::10
Nexthop local: fe80::5054:ff:fe11:d4f5
BGP connection: shared network
Last Reset: 00:11:06, due to OPEN Message Error (Notification sent)
Notification Error Message: (OPEN Message Error/Bad Peer AS.)
```

R2

```
R1#show bgp ipv6 neighbors
BGP neighbor is 3ffe:10::11, remote AS 11, local AS 10, external link
  BGP version 4, local router ID 192.168.10.10, remote router ID 192.168.10.11
  BGP state = Established, up for 00:11:02
  Last read 00:00:23, 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 140 messages, 0 notifications, 0 in queue
  Sent 194 messages, 113 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv6 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Inbound path policy configured
  Route map for incoming advertisements is *myPrefixList
```

```
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
Local host: 3ffe:10::10, Local port: 179
Foreign host: 3ffe:10::11, Foreign port: 46796
Nexthop: 192.168.10.10
Nexthop global: 3ffe:10::10
Nexthop local: fe80::5054:ff:fe11:d4f5
BGP connection: shared network
Last Reset: 00:11:06, due to OPEN Message Error (Notification sent)
Notification Error Message: (OPEN Message Error/Bad Peer AS.)
```

Route Reflector

Use Route Reflectors to reduce the iBGP mesh inside an AS. 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 RR only, and 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 the R2. This reduces the iBGP peer connections in AS1.

Topology

Figure 70. BGP4+ Route Reflector

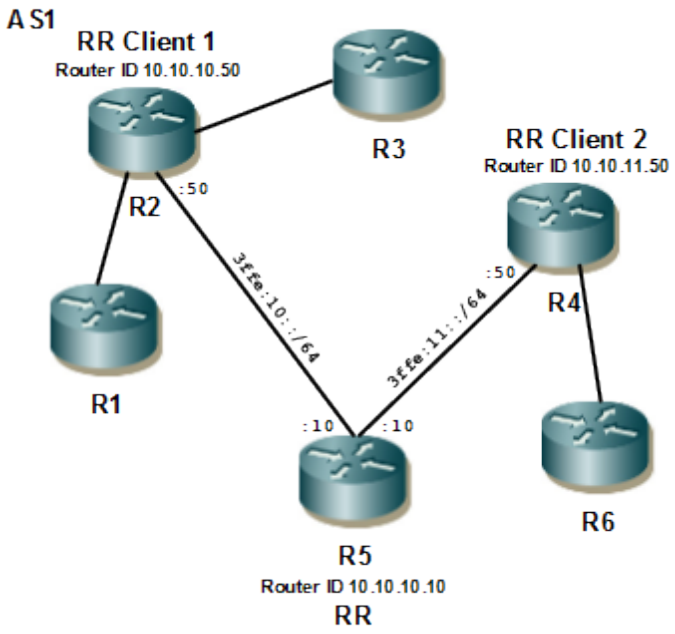

Configuration

RR (R5)

`#configure terminal`	Enter Configure mode.
`(config)#router bgp 1`	Define the routing process. The number `1`

	specifies the AS number of R5 (RR).
(config-router)#bgp router-id 10.10.10.10	Configure a fixed Router ID (10.10.10.10) for the BGP4+ routing process.
(config-router)#neighbor 3ffe:10::50 remote-as 1	Define the BGP neighbor (R2), and establish a TCP session by specifying the IPv6 address (3ffe:10::50) and the AS number (1) of neighbor R2.
(config-router)#neighbor 3ffe:11::50 remote-as 1	Define the BGP neighbor (R4), and establish a TCP session by specifying the IPv6 address (3ffe:11::50) and the AS number (1) of neighbor R4.
(config-router)#address-family ipv6	Enter Address Family mode for configuring routing sessions that use IPv6 address prefixes.
(config-router-af)#neighbor 3ffe:10::50 activate	Activate the neighbor R2 (3ffe:10::50), and enable exchange of IPv6 address prefix types with this neighbor.
(config-router-af)#neighbor 3ffe:10::50 route-reflector-client	Configure R5 as the Route-Reflector (RR) and neighbor R2 as its client.
(config-router-af)#neighbor 3ffe:11::50 activate	Activate the neighbor R4 (3ffe:11::50), and enable exchange of IPv6 address prefix types with this neighbor.
(config-router-af)#neighbor 3ffe: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
(config-router)#commit	Commit the candidate configuration to the running configuration.

RR Client 1 (R2)

(config)#router bgp 1	Define the routing process. The number 1 specifies the AS number of R2 (RR Client 1).
(config-router)#bgp router-id 10.10.10.50	Configure a fixed Router ID (10.10.10.50) for the BGP4+ routing process.
(config-router)#neighbor 3ffe:10::10 remote-as 1	Define the BGP neighbor (R5), and establish a TCP session by specifying the IPv6 address (3ffe:10::10) and the AS number (1) of neighbor R5.
(config-router)#address-family ipv6	Enter Address Family mode for configuring routing sessions that use IPv6 address prefixes.
(config-router-af)#neighbor 3ffe:10::10 activate	Activate the neighbor (3ffe:10::10), and enable exchange of IPv6 address prefix types with this neighbor.
(config-router-af)#exit-address-family	Exit address family

`(config-router)#commit`	Commit the candidate configuration to the running configuration.

RR Client 2 (R4)

`(config)#router bgp 1`	Define the routing process. The number `1` specifies the AS number of `R4` (RR Client 2).
`(config-router)#bgp router-id 10.10.11.50`	Configure a fixed Router ID (`10.10.11.50`) for the BGP4+ routing process.
`(config-router)#neighbor 3ffe:11::10 remote-as 1`	Define the BGP neighbor (`R5`), and establish a TCP session by specifying the IPv6 address (`3ffe:11::10`) and the AS number `(1)` of the neighbor.
`(config-router)#address-family ipv6`	Enter Address Family mode for configuring routing sessions that use IPv6 address prefixes.
`(config-router-af)#neighbor 3ffe:11::10 activate`	Activate the neighbor `(3ffe:11::10)`, and enable exchange of IPv6 address prefix types with this neighbor.
`(config-router-af)#exit-address-family`	Exit address family
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

show bgp ipv6 summary

```
1#show bgp ipv6 summary
BGP router identifier 223.223.223.1, local AS number 6500001
BGP table version is 1
3 BGP AS-PATH entries
0 BGP community entries

Neighbor           V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
10::13             4 6500001          0          0       0      0      0     00:13:17
          EST

Total number of neighbors 1

Total number of Established sessions  1
R1#
```

show bgp ipv6 neighbors

```
show ip bgp neighbors
--------------------------------
R1#sh ip bgp neighbors
BGP neighbor is 10.1.11.1, remote AS 99, local AS 6500001, external link, peer index: 3
  BGP version 4, local router ID 223.223.223.1, remote router ID 192.0.0.3
  BGP state = Established, up for 05:39:44, Socket state = Connected
  Last read 00:00:14, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised
    4-Octet ASN Capability: advertised
    Address family IPv4 Unicast: advertised and received
```

```
  Received 785 messages, 0 notifications, 0 in queue
  Sent 805 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)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  3 announced prefixes

 Connections established 1; dropped 0
Local host: 10.1.11.0, Local port: 62958
Foreign host: 10.1.11.1, Foreign port: 179
TCP MSS: (0), Advertise TCP MSS: (1460), Send TCP MSS: (1460),  Receive TCP MSS: (1460)
Sock FD : (32)
Nexthop: 10.1.11.0 cd25/1
Nexthop global: :: cd25/1
Nexthop local: :: cd25/1
BGP connection: non shared network

BGP neighbor is 10.1.12.1, remote AS 6500001, local AS 6500001, internal link, peer index: 4
  BGP version 4, local router ID 223.223.223.1, remote router ID 0.0.0.0
  BGP state = Idle, Socket state = Idle
  Last read Never, hold time is 90, keepalive interval is 30 seconds
  Received 0 messages, 0 notifications, 0 in queue
  Sent 0 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 0
  Index 2, Offset 0, Mask 0x4
  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 0; dropped 0
TCP MSS: (0), Advertise TCP MSS: (0), Send TCP MSS: (536),  Receive TCP MSS: (0)
Sock FD : (36)


BGP neighbor is 10.1.13.1, remote AS 6500001, local AS 6500001, internal link, peer index: 5
  BGP version 4, local router ID 223.223.223.1, remote router ID 0.0.0.0
  BGP state = Idle, Socket state = Idle
  Last read Never, hold time is 90, keepalive interval is 30 seconds
  Received 126 messages, 1 notifications, 0 in queue
  Sent 119 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 0
  Index 3, Offset 0, Mask 0x8
  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 1; dropped 1
TCP MSS: (0), Advertise TCP MSS: (0), Send TCP MSS: (536),  Receive TCP MSS: (0)
Sock FD : (37)
Last Reset: 04:49:44, due to BGP Notification received
Notification Error Message: (Cease/Peer Unconfigured.)


R1#
```

Confederations

In this example, AS1 contains three Confederated Autonomous Systems--AS 1000, AS 1001 and AS 1002. To any outside AS, the overall Confederation is a single AS, AS1. Confederation eBGP is run between R2 and R5, and between R5 and R7. R2 is configured so that its local AS is 1000. Its peer connection to R5 is set up like any other eBGP session. The `bgp confederation identifier` command informs the router that it is a member of a Confederation and passes the Confederation ID. The `bgp confederation peers` command lists the member AS to which R2 is connected. The same command tells the BGP process that the eBGP connection is a Confederation eBGP, rather than a normal eBGP.

Topology

Figure 71. BGP4+ Confederations

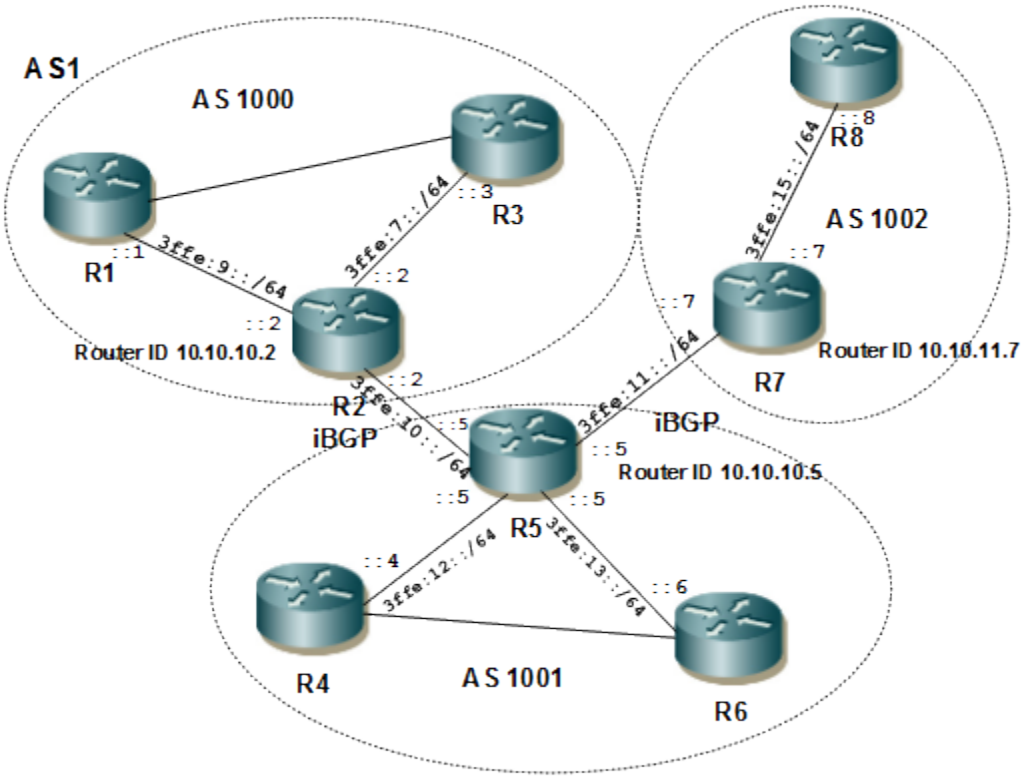

Configuration

R2

`#configure terminal`	Enter Configure mode.
`(config)#router bgp 1000`	Define the routing process. The number `1000` specifies the AS number of `R2`.
`(config-router)#bgp router-id 10.10.10.2`	Configure a fixed Router ID (`10.10.10.2`) for the BGP4+ routing process.
`(config-router)#bgp confederation identifier 1`	Specify the BGP Confederation Identifier (1). To others, the group will appear as a single AS, and the

	identifier as its AS number.
(config-router)#bgp confederation peers 1001	Specify AS 1001 as s confederation peer, making it a member of the Confederation.
(config-router)#bgp confederation peers 1002	Specify AS 1002 as s confederation peer, making it a member of the Confederation.
(config-router)#neighbor 3ffe:10::5 remote-as 1001	Define the BGP neighbor (R5), and establish a TCP session by specifying the IPv6 address (3ffe:10::5) and the AS number (1001) of neighbor R5.
(config-router)#neighbor 3ffe:9::1 remote-as 1000	Define the BGP neighbor (R1), and establish a TCP session by specifying the IPv6 address (3ffe:9::1) and the AS number (1000) of neighbor R1.
(config-router)#neighbor 3ffe:7::3 remote-as 1000	Define the BGP neighbor (R3), and establish a TCP session by specifying the IPv6 address (3ffe:7::3) and the AS number (1000) of neighbor R3.
(config-router)#address-family ipv6	Enter Address Family mode for configuring routing sessions that use IPv6 address prefixes.
(config-router-af)#neighbor 3ffe:10::5 activate	Activate neighbor R5(3ffe:10::5), and enable exchange of IPv6 address prefix types with this neighbor.
(config-router-af)#neighbor 3ffe:9::1 activate	Activate neighbor R1(3ffe:9::1), and enable exchange of IPv6 address prefix types with this neighbor.
(config-router-af)#neighbor 3ffe:7::3 activate	Activate neighbor R3(3ffe:7::3), and enable exchange of IPv6 address prefix types with this neighbor.
(config-router-af)#exit-address-family	Exit address family
(config-router)#commit	Commit the candidate configuration to the running configuration.

R5

(config)#router bgp 1001	Define the routing process. The number 1001 specifies the AS number of R5.
(config-router)#bgp router-id 10.10.10.5	Configure a fixed Router ID (10.10.10.5) for the BGP4+ routing process.
(config-router)#bgp confederation identifier 1	Specify the BGP Confederation Identifier (1). To others, the group will appear as a single AS, and the identifier as its AS number.
(config-router)#bgp confederation peers 1000	Specify AS 1000 as s confederation peer, making it a member of the Confederation.
(config-router)#bgp confederation peers 1002	Specify AS 1002 as s confederation peer, making it a

	member of the Confederation.
`(config-router)#neighbor 3ffe:10::2 remote-as 1000`	Define the BGP neighbor (`R2`), and establish a TCP session by specifying the IPv6 address (`3ffe:10::2`) and the AS number (`1000`) of neighbor `R2`.
`(config-router)#neighbor 3ffe:11::7 remote-as 1002`	Define the BGP neighbor (`R7`), and establish a TCP session by specifying the IPv6 address (`3ffe:11::7`) and the AS number (`1002`) of neighbor `R7`.
`(config-router)#neighbor 3ffe:12::4 remote-as 1001`	Define the BGP neighbor (`R4`), and establish a TCP session by specifying the IPv6 address (`3ffe:12::4`) and the AS number (`1001`) of neighbor `R4`.
`(config-router)#neighbor 3ffe:13::6 remote-as 1001`	Define the BGP neighbor (`R6`), and establish a TCP session by specifying the IPv6 address (`3ffe:13::6`) and the AS number (`1001`) of neighbor `R6`.
`(config-router)#address-family ipv6`	Enter Address Family mode for configuring routing sessions that use IPv6 address prefixes.
`(config-router-af)#neighbor 3ffe:10::2 activate`	Activate the neighbor `R2(3ffe:10::2)`, and enable exchange of IPv6 address prefix types with this neighbor.
`(config-router-af)#neighbor 3ffe:11::7 activate`	Activate the neighbor `R7(3ffe:11::7)`, and enable exchange of IPv6 address prefix types with this neighbor.
`(config-router-af)#neighbor 3ffe:12::4 activate`	Activate the neighbor `R4(3ffe:12::4)`, and enable exchange of IPv6 address prefix types with this neighbor.
`(config-router-af)#neighbor 3ffe:13::6 activate`	Activate the neighbor `R6(3ffe:13::6)`, and enable exchange of IPv6 address prefix types with this neighbor.
`(config-router-af)#exit-address-family`	Exit address family
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

R7

`(config)#router bgp 1002`	Define the routing process. The number `1002` specifies the AS number of `R7`.
`(config-router)#bgp router-id 10.10.11.7`	Configure a fixed Router ID (`10.10.11.7`) for the BGP4+ routing process.
`(config-router)#bgp confederation identifier 1`	Specify BGP Confederation Identifier (`1`). To others, the group will appear as a single AS, and the identifier as its AS number.

© 2026 IP Infusion Inc. Proprietary

`(config-router)#bgp confederation peers 1000`	Specify AS `1000` as s confederation peer, making it a member of the Confederation.
`(config-router)#bgp confederation peers 1001`	Specify AS `1001` as s confederation peer, making it a member of the Confederation.
`(config-router)#neighbor 3ffe:11::5 remote-as 1001`	Define the BGP neighbor (`R5`), and establish a TCP session by specifying the IPv6 address (`3ffe:11::5`) and the AS number `(1001)` of neighbor R5.
`(config-router)#neighbor 3ffe:15::8 remote-as 1002`	Define the BGP neighbor (`R8`), and establish a TCP session by specifying the IPv6 address (`3ffe:15::8`) and the AS number `(1002)` of neighbor `R8`.
`(config-router)#address-family ipv6`	Enter Address Family mode for configuring routing sessions that use IPv6 address prefixes.
`(config-router-af)#neighbor 3ffe:11::5 activate`	Activate the neighbor `R5(3ffe:11::5)`, and enable exchange of IPv6 address prefix types with this neighbor.
`(config-router-af)#neighbor 3ffe:15::8 activate`	Activate the neighbor `R8(3ffe:15::8)`, and enable exchange of IPv6 address prefix types with this neighbor.
`(config-router-af)#exit-address-family`	Exit address family
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

show bgp ipv6 summary

```
1#show bgp ipv6 summary
BGP router identifier 223.223.223.1, local AS number 6500001
BGP table version is 1
3 BGP AS-PATH entries
0 BGP community entries

Neighbor           V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
10::13             4 6500001          0          0       0      0      0     00:13:17
          EST

Total number of neighbors 1

Total number of Established sessions  1
R1#
```

show bgp ipv6 neighbors

```
show ip bgp neighbors
---------------------------------
R1#sh ip bgp neighbors
BGP neighbor is 10.1.11.1, remote AS 99, local AS 6500001, external link, peer index: 3
  BGP version 4, local router ID 223.223.223.1, remote router ID 192.0.0.3
  BGP state = Established, up for 05:39:44, Socket state = Connected
  Last read 00:00:14, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
```

```
    Route refresh: advertised
    4-Octet ASN Capability: advertised
    Address family IPv4 Unicast: advertised and received
  Received 785 messages, 0 notifications, 0 in queue
  Sent 805 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)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  3 announced prefixes

 Connections established 1; dropped 0
Local host: 10.1.11.0, Local port: 62958
Foreign host: 10.1.11.1, Foreign port: 179
TCP MSS: (0), Advertise TCP MSS: (1460), Send TCP MSS: (1460),  Receive TCP MSS: (1460)
Sock FD : (32)
Nexthop: 10.1.11.0 cd25/1
Nexthop global: :: cd25/1
Nexthop local: :: cd25/1
BGP connection: non shared network

BGP neighbor is 10.1.12.1, remote AS 6500001, local AS 6500001, internal link, peer index: 4
  BGP version 4, local router ID 223.223.223.1, remote router ID 0.0.0.0
  BGP state = Idle, Socket state = Idle
  Last read Never, hold time is 90, keepalive interval is 30 seconds
  Received 0 messages, 0 notifications, 0 in queue
  Sent 0 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 0
  Index 2, Offset 0, Mask 0x4
  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 0; dropped 0
TCP MSS: (0), Advertise TCP MSS: (0), Send TCP MSS: (536),  Receive TCP MSS: (0)
Sock FD : (36)


BGP neighbor is 10.1.13.1, remote AS 6500001, local AS 6500001, internal link, peer index: 5
  BGP version 4, local router ID 223.223.223.1, remote router ID 0.0.0.0
  BGP state = Idle, Socket state = Idle
  Last read Never, hold time is 90, keepalive interval is 30 seconds
  Received 126 messages, 1 notifications, 0 in queue
  Sent 119 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 0
  Index 3, Offset 0, Mask 0x8
  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 1; dropped 1
TCP MSS: (0), Advertise TCP MSS: (0), Send TCP MSS: (536),  Receive TCP MSS: (0)
Sock FD : (37)
Last Reset: 04:49:44, due to BGP Notification received
```

```
Notification Error Message: (Cease/Peer Unconfigured.)


R1#
```

BGP4+ Graceful Restart

Using BGP+ graceful restart, the data-forwarding plane of a router can continue to process and forward packets, even if the control plane (which is responsible for determining best paths) fails.

Topology

Figure 72. BGP4+ Graceful Restart

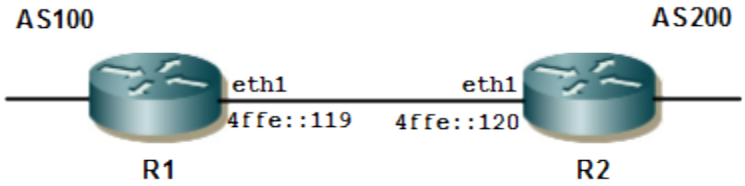

Configuration

R1

`#configure terminal` `(config)#interface lo`	Enter Configure mode. Enter to the interface mode.
`(config-if)#ipv6 address 4000::1/64`	Assign ipv6 address to loopback interface.
`(config-if)#exit`	Exit from interface mode.
`(config)#interface xe9`	Enter to the interface mode.
`(config-if)#ipv6 address 4ffe::119/64`	Configure ipv6 address.
`(config-if)#exit`	Exit from interface mode.
`(config)#router bgp 100`	Assign the ASN value (100) to the router.
`(config-router)# bgp router-id 1.1.1.1`	Configuring bgp router id.
`(config-router)#bgp graceful-restart`	Enable BGP graceful restart support.
`(config-router)#neighbor 4ffe::120 remote-as 200`	Specify the neighbor's IP address (4ffe::120) and the ASN value of the neighbor (200).
`(config-router)#address-family ipv6 unicast`	Exchange the IPv6 capabilities, and switch the mode to the IPv6 address family.
`(config-router-af)#network 4000::/64`	Advertising ipv6 address.
`(config-router-af)#neighbor 4ffe::120 activate`	Specify the neighbor's IPv6 address (4ffe::120) , and activate the neighbor.
`(config-router-af)#neighbor 4ffe::120 capability graceful-restart`	Specify the neighbor's IPv6 address (4ffe::120) for which the graceful restart capability is supported.

#configure terminal	Enter Configure mode.
(config-router-af)#exit-address-family	Exit address family
(config-router)#commit	Commit the candidate configuration to the running configuration.

R2

#configure terminal	Enter Configure mode
(config)#interface xe9	Enter to the interface mode
(config-if)#ipv6 address 4ffe::119/64	Configure ipv6 address.
(config-if)#exit	Exit from interface mode.
(config)#router bgp 200	Assign the ASN value (200) to the router.
(config-router)# bgp router-id 2.2.2.2	Configuring bgp router id.
(config-router)#bgp graceful-restart	Enable BGP graceful restart support.
(config-router)#bgp graceful-restart restart-time 120	Configure the maximum time (120) required for neighbor(s) to restart.
(config-router)#bgp graceful-restart stalepath-time 120	Configure the maximum time (120) to retain stale paths from the restarting neighbor(s).
(config-router)#neighbor 4ffe::119 remote-as 100	Specify the neighbor's IP address (4ffe::119) and the ASN value of the neighbor (100).
(config-router)#address-family ipv6 unicast	Exchange the IPv6 capabilities, and switch the mode to the IPv6 address family.
(config-router-af)#neighbor 4ffe::119 activate	Specify the neighbor's IPv6 address（4ffe::119）, and activate the neighbor.
(config-router-af)#neighbor 4ffe::119 capability graceful-restart	Specify the neighbor's IPv6 address (4ffe::119) for which the graceful restart capability is supported.
(config-router-af)#exit-address-family	Exit address family
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

R1

```
#show bgp ipv6 summary
BGP router identifier 1.1.1.1, 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
4ffe::120                4   200    8         11       1      0      0  00:01:02               0

Total number of neighbors 1

Total number of Established sessions 1
```

```
#show ip bgp neighbors
BGP neighbor is 4ffe::120, remote AS 200, 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 00:01:31
  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 IPv6 Unicast: advertised and received
  Received 9 messages, 0 notifications, 0 in queue
  Sent 12 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv6 Unicast
  BGP table version 1, neighbor version 1
  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)
  0 accepted prefixes
  1 announced prefixes

 Connections established 2; dropped 1
 Graceful-restart Status:
  Remote restart-time is 90 sec

Local host: 4ffe::119, Local port: 179
Foreign host: 4ffe::120, Foreign port: 58264
Nexthop: 1.1.1.1
Nexthop global: 4ffe::119
Nexthop local: fe80::eac5:7aff:fe8b:a827
BGP connection: shared network

#show bgp ipv6
BGP table version is 1, local router ID is 1.1.1.1
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
*>   4000::/64        ::                             0        100       32768  i

Total number of prefixes 1

#show ipv6 route database bgp
IP Route Table for VRF "default"

#show ipv6 route database
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
       > - selected route, * - FIB route,p - stale info
Timers: Uptime

IP Route Table for VRF "default"
C   *>  ::1/128 via ::, lo, 10:38:02
C   *>  4000::/64 via ::, lo, 00:25:40
C   *>  4ffe::/64 via ::, xe9, 00:49:30
C   *>  fe80::/64 via ::, xe9, 02:06:22
C       fe80::/64 via ::, xe5, 05:16:16
C       fe80::/64 via ::, xe12, 05:33:37
C       fe80::/64 via ::, xe11, 05:33:37
C       fe80::/64 via ::, ce1, 10:38:02
```

© 2026 IP Infusion Inc. Proprietary

```
#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, 10:38:34
C       4000::/64 via ::, lo, 00:26:12
C       4ffe::/64 via ::, xe9, 00:50:02
C       fe80::/64 via ::, xe9, 02:06:54
```

R2

```
#show bgp ipv6 summary
BGP router identifier 2.2.2.2, local AS number 200
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
4ffe::119                4   100   14         14       2      0      0  00:05:08               1

Total number of neighbors 1

Total number of Established sessions 1
#show ip bgp neighbors
BGP neighbor is 4ffe::119, remote AS 100, local AS 200, external link
  BGP version 4, local router ID 2.2.2.2, remote router ID 1.1.1.1
  BGP state = Established, up for 00:06:16
  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 IPv6 Unicast: advertised and received
  Received 17 messages, 0 notifications, 0 in queue
  Sent 17 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv6 Unicast
  BGP table version 2, neighbor version 2
  Index 1, Offset 0, Mask 0x2
  AF-dependant capabilities:
    Graceful restart: advertised, received

  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 0
 Graceful-restart Status:
  Remote restart-time is 90 sec

Local host: 4ffe::120, Local port: 58264
Foreign host: 4ffe::119, Foreign port: 179
Nexthop: 2.2.2.2
Nexthop global: 4ffe::120
Nexthop local: fe80::82a2:35ff:fe7e:8833
BGP connection: shared network

#show bgp ipv6
BGP table version is 2, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,
```

© 2026 IP Infusion Inc. Proprietary

```
             l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop                    Metric    LocPrf     Weight Path
*>   4000::/64        4ffe::119(fe80::eac5:7aff:fe8b:a827)
                                                    0        100       0   100 i

Total number of prefixes 1

#show ipv6 route database bgp
IP Route Table for VRF "default"
B   *>  4000::/64 [20/0] via fe80::eac5:7aff:fe8b:a827, xe9, 00:06:53


#show ipv6 route database
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
       > - selected route, * - FIB route,p - stale info
Timers: Uptime

IP Route Table for VRF "default"
C   *>  ::1/128 via ::, lo, 02:11:57
B   *>  4000::/64 [20/0] via fe80::eac5:7aff:fe8b:a827, xe9, 00:08:19
C   *>  4ffe::/64 via ::, xe9, 00:50:24
C   *>  fe80::/64 via ::, xe13, 00:05:07
C       fe80::/64 via ::, ce1, 02:11:57
C       fe80::/64 via ::, ce0, 02:11:57
C       fe80::/64 via ::, xe23, 02:11:57
C       fe80::/64 via ::, xe22, 02:11:57
C       fe80::/64 via ::, xe21, 02:11:57
C       fe80::/64 via ::, xe20, 02:11:57
C       fe80::/64 via ::, xe15, 02:11:57
C       fe80::/64 via ::, xe9, 02:11:57
C       fe80::/64 via ::, xe6, 02:11:57

#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, 02:14:13
B       4000::/64 [20/0] via fe80::eac5:7aff:fe8b:a827, xe9, 00:10:35
C       4ffe::/64 via ::, xe9, 00:52:40
C       fe80::/64 via ::, xe13, 00:07:23
```

Performing BGP Graceful Restart

```
#write memory
Building configuration...
[OK]
#restart bgp graceful
%Warning : BGP process will stop and needs to restart manually,
You may lose bgp configuration,if not saved
Proceed for graceful restart? (y/n):y
2019 Feb 22 16:08:20.149 : OcNOS : BGP : CRITI : [BGP_OPR_NEIGH_STATE_DOWN_2]: Neighbour [4ffe::119]
Session down as GR configured/unconfigured
2019 Feb 22 16:08:20.149 : OcNOS : BGP : CRITI : [BGP_OPR_NEIGH_STATE_DOWN_2]: Neighbour [4ffe::119]
Session down due to config deletion
```

```
2019 Feb 22 16:08:20.153 : OcNOS : CML : CRITI : Module bgpd disconnected with CML
%% Managed module is down or crashed
2019 Feb 22 16:08:20.153 : OcNOS : CML : CRITI : Module bgpd disconnected with CML

#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, 02:18:54
B       4000::/64 [20/0] via fe80::eac5:7aff:fe8b:a827, xe9, 00:15:16
C       4ffe::/64 via ::, xe9, 00:57:21
C       fe80::/64 via ::, xe13, 00:12:04

#show ipv6 route database
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
       > - selected route, * - FIB route,p - stale info
Timers: Uptime

IP Route Table for VRF "default"
C   *>  ::1/128 via ::, lo, 02:19:23
B   *>p 4000::/64 [20/0] via fe80::eac5:7aff:fe8b:a827, xe9, 00:15:45
C   *>  4ffe::/64 via ::, xe9, 00:57:50
C   *>  fe80::/64 via ::, xe13, 00:12:33
C       fe80::/64 via ::, ce1, 02:19:23
C       fe80::/64 via ::, ce0, 02:19:23
C       fe80::/64 via ::, xe23, 02:19:23
C       fe80::/64 via ::, xe22, 02:19:23
C       fe80::/64 via ::, xe21, 02:19:23
C       fe80::/64 via ::, xe20, 02:19:23
C       fe80::/64 via ::, xe15, 02:19:23
C       fe80::/64 via ::, xe9, 02:19:23
C       fe80::/64 via ::, xe6, 02:19:23
```

Restarting the BGP Daemon

```
#start-shell
bash-5.0$ su
Password:
root@OcNOS:/home/ocnos# cd /usr/local/sbin/
root@OcNOS:/usr/local/sbin# ./bgpd -d
```

Configure BGP4+ Distance

Administrative distance in BGP+ can be configured for a specific address family.

This example shows configuring the BGP administrative distance for the IPv6 address family.

Topology

Figure 73. BGP4+ Distance

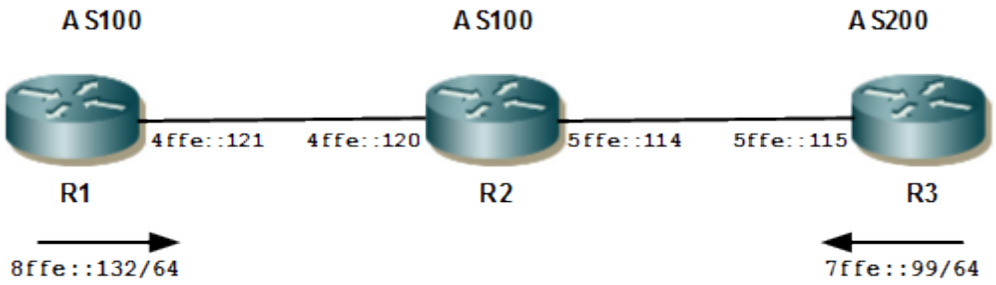

Configuration

R1

`#configure terminal`	Enter Configure mode.
`(config)#router bgp 100`	Assign the ASN value (100) to the router.
`(config-router)#neighbor 4ffe::120 remote-as 100`	Specify the neighbor's IP address and ASN value.
`(config-router)#address-family ipv6 unicast`	Enter IPv6 Address-Family mode.
`(config-router-af)#neighbor 4ffe::120 activate`	Activate the IPv6 neighbor.
`(config-router-af)#network 8ffe::132/64`	Specify the network to be advertised by the BGP routing process.
`(config-router-af)#exit-address-family`	Exit Address Family mode, and return to Router mode.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

R2

`#configure terminal`	Enter Configure mode.
`(config)#router bgp 100`	Assign the ASN value (100) to the router.
`(config-router)#neighbor 4ffe::121 remote-as 100`	Specify the neighbor’s IP address and ASN value.
`(config-router)#neighbor 5ffe::115 remote-as 200`	Specify the neighbor’s IP address and the ASN value of another neighbor.
`(config-router)#address-family ipv6 unicast`	Enter IPv6 Address-Family mode.
`(config-router-af)#aggregate-address 2ffe::102/32 summary-only`	Configure an IPv6 non-AS-set aggregate route on R2. The local distance will be applied to this route.
`(config-router-af)#neighbor 4ffe::121 activate`	Activate the IPv6 neighbor.
`(config-router-af)#neighbor 5ffe::115 activate`	Activate the IPv6 neighbor.
`(config-router-af)#distance bgp 12 13 11`	Configure the administrative distance for external, internal, and local routes received in IPv6 Address-

	Family mode.
(config-router-af)#exit-address-family	Exit Address Family mode, and return to Router mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

R3

#configure terminal	Enter Configure mode.
(config)#router bgp 100	Assign the ASN value (100) to the router.
(config-router)#neighbor 5ffe::114 remote-as 100	Specify the neighbor's IP address and ASN value.
(config-router)#address-family ipv6 unicast	Enter IPv6 Address-Family mode.
(config-router)#network 7ffe::99/64	Specify the network to be advertised by the BGP routing process.
(config-router-af)#neighbor 5ffe::114 activate	Activate the IPv6 neighbor.
(config-router-af)#exit-address-family	Exit Address Family mode, and return to Router mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

show bgp ipv6 summary

```
1#show bgp ipv6 summary
BGP router identifier 223.223.223.1, local AS number 6500001
BGP table version is 1
3 BGP AS-PATH entries
0 BGP community entries

Neighbor           V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
10::13             4 6500001          0          0       0      0      0     00:13:17
          EST

Total number of neighbors 1

Total number of Established sessions  1
R1#
```

show ip bgp neighbors

```
show ip bgp neighbors
---------------------------------
R1#sh ip bgp neighbors
BGP neighbor is 10.1.11.1, remote AS 99, local AS 6500001, external link, peer index: 3
  BGP version 4, local router ID 223.223.223.1, remote router ID 192.0.0.3
  BGP state = Established, up for 05:39:44, Socket state = Connected
  Last read 00:00:14, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised
    4-Octet ASN Capability: advertised
    Address family IPv4 Unicast: advertised and received
```

```
  Received 785 messages, 0 notifications, 0 in queue
  Sent 805 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)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  3 announced prefixes

 Connections established 1; dropped 0
Local host: 10.1.11.0, Local port: 62958
Foreign host: 10.1.11.1, Foreign port: 179
TCP MSS: (0), Advertise TCP MSS: (1460), Send TCP MSS: (1460),  Receive TCP MSS: (1460)
Sock FD : (32)
Nexthop: 10.1.11.0 cd25/1
Nexthop global: :: cd25/1
Nexthop local: :: cd25/1
BGP connection: non shared network

BGP neighbor is 10.1.12.1, remote AS 6500001, local AS 6500001, internal link, peer index: 4
  BGP version 4, local router ID 223.223.223.1, remote router ID 0.0.0.0
  BGP state = Idle, Socket state = Idle
  Last read Never, hold time is 90, keepalive interval is 30 seconds
  Received 0 messages, 0 notifications, 0 in queue
  Sent 0 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 0
  Index 2, Offset 0, Mask 0x4
  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 0; dropped 0
TCP MSS: (0), Advertise TCP MSS: (0), Send TCP MSS: (536),  Receive TCP MSS: (0)
Sock FD : (36)


BGP neighbor is 10.1.13.1, remote AS 6500001, local AS 6500001, internal link, peer index: 5
  BGP version 4, local router ID 223.223.223.1, remote router ID 0.0.0.0
  BGP state = Idle, Socket state = Idle
  Last read Never, hold time is 90, keepalive interval is 30 seconds
  Received 126 messages, 1 notifications, 0 in queue
  Sent 119 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 0
  Index 3, Offset 0, Mask 0x8
  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 1; dropped 1
TCP MSS: (0), Advertise TCP MSS: (0), Send TCP MSS: (536),  Receive TCP MSS: (0)
Sock FD : (37)
Last Reset: 04:49:44, due to BGP Notification received
Notification Error Message: (Cease/Peer Unconfigured.)


R1#
```

show ipv6 route database bgp

```
R1#show ipv6 route database bgp
IP Route Table for VRF "default"
B   *>  10:10::/64 [200/0] via fe80::5e07:58ff:fe54:28a2, cd24/4, installed 00:01:27, last update
00:01:27 ago

Total number of IPv6 routes 1
R1#
```

show ipv6 route database

```
R1#show ipv6 route database
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
       > - selected route, * - FIB route,p - stale info
Timers: Uptime

IP Route Table for VRF "default"
C   *>  ::1/128 via ::, lo, installed 00:24:11, last update 00:24:11 ago
C   *>  10::/64 via ::, cd24/1, installed 00:03:19, last update 00:03:19 ago
C   *>  fe80::/64 via ::, cd24/1, installed 00:03:33, last update 00:03:33 ago
C       fe80::/64 via ::, cd25/3, installed 00:23:49, last update 00:23:49 ago
C       fe80::/64 via ::, cd24/4, installed 00:23:49, last update 00:23:49 ago
C       fe80::/64 via ::, cd25/1, installed 00:23:49, last update 00:23:49 ago

Total number of IPv6 routes 6
R1#
```

show ipv6 route

```
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,
 P - SRV6-POLICY,
        v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 00:23:49, last update 00:23:49 ago
C       10::/64 via ::, cd24/1, installed 00:02:57, last update 00:02:57 ago
C       fe80::/64 via ::, cd24/1, installed 00:03:11, last update 00:03:11 ago
R1#
```

show bgp ipv6

```
R1#show bgp ipv6
BGP table version is 2, local router ID is 223.223.223.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
              l - labeled, S Stale, x-EVPN, w-Stitched-routes
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

    Network          Next Hop                    Metric    LocPrf     Weight Path
*>    10::/64          ::                             0        100       32768  ?        -
* i                    10::13(fe80::eac5:7aff:feb2:d1d)
                                                     0        100       0    ?        -

Total number of prefixes 1
R1#
```

BGP4+ Graceful Reset

The graceful restart mechanism for BGP+ session reset (the BGP+ daemon is not restarted) is used so that any changes in network configuration do not affect packet forwarding. The bgp graceful-restart graceful-reset CLI invokes graceful restart when a configuration change forces a peer reset. Graceful restart is invoked only when these CLI configuration changes force a peer reset.

Topology

Figure 74. BGP4+ Graceful Reset

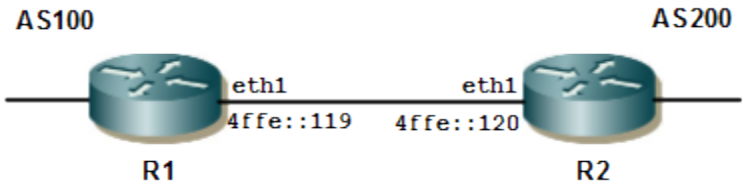

Configuration

R1

`#configure terminal` `(config)#router bgp 100`	Enter Configure mode. Assign the ASN value (100) to the router.
`(config-router)#bgp graceful-restart`	Enable BGP graceful restart support.
`(config-router)#bgp graceful-restart graceful-reset`	Configure to invoke graceful restart when a configuration change forces a peer reset.
`(config-router)#neighbor 4ffe::120 remote-as 200`	Specify the neighbor's IP address (4ffe::120) and the ASN value of the neighbor (200).
`(config-router)#address-family ipv6 unicast`	Exchange the IPv6 capabilities, and switch the mode to the IPv6 address family.
`(config-router-af)#neighbor 4ffe::120 activate`	Specify the neighbor's IPv6 address (4ffe::120) , and activate the neighbor.
`(config-router-af)#neighbor 4ffe::120 capability graceful-restart`	Specify the neighbor's IPv6 address (4ffe::120) for which the graceful restart capability is supported.
`(config-router-af)#exit-address-family`	Exit Address Family mode, and return to Router mode.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

R2

`#configure terminal` `(config)#router bgp 200`	Enter Configure mode. Assign the ASN value (200) to the router.
`(config-router)#bgp graceful-restart`	Enable BGP graceful restart support.

© 2026 IP Infusion Inc. Proprietary

`#configure terminal`	Enter Configure mode.
`(config-router)#bgp graceful-restart graceful-reset`	Configure to invoke graceful restart when a configuration change forces a peer reset.
`(config-router)#neighbor 4ffe::119 remote-as 100`	Specify the neighbor's IP address (4ffe::119) and the ASN value of the neighbor (100).
`(config-router)#address-family ipv6 unicast`	Exchange the IPv6 capabilities, and switch the mode to the IPv6 address family.
`(config-router-af)#neighbor 4ffe::119 activate`	Specify the neighbor's IPv6 address (4ffe::119) , and activate the neighbor.
`(config-router-af)#neighbor 4ffe::119 capability graceful-restart`	Specify the neighbor's IPv6 address (4ffe::119) for which the graceful restart capability is supported.
`(config-router-af)#exit-address-family`	Exit Address Family mode, and return to Router mode.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

show bgp ipv6 summary

```
1#show bgp ipv6 summary
BGP router identifier 223.223.223.1, local AS number 6500001
BGP table version is 1
3 BGP AS-PATH entries
0 BGP community entries

Neighbor           V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/PfxRcd   Desc
10::13             4 6500001          0          0       0      0      0     00:13:17
          EST

Total number of neighbors 1

Total number of Established sessions  1
R1#
```

show ip bgp neighbors

```
show ip bgp neighbors
---------------------------------
R1#sh ip bgp neighbors
BGP neighbor is 10.1.11.1, remote AS 99, local AS 6500001, external link, peer index: 3
  BGP version 4, local router ID 223.223.223.1, remote router ID 192.0.0.3
  BGP state = Established, up for 05:39:44, Socket state = Connected
  Last read 00:00:14, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised
    4-Octet ASN Capability: advertised
    Address family IPv4 Unicast: advertised and received
  Received 785 messages, 0 notifications, 0 in queue
  Sent 805 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)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
```

```
  3 announced prefixes

 Connections established 1; dropped 0
Local host: 10.1.11.0, Local port: 62958
Foreign host: 10.1.11.1, Foreign port: 179
TCP MSS: (0), Advertise TCP MSS: (1460), Send TCP MSS: (1460),  Receive TCP MSS: (1460)
Sock FD : (32)
Nexthop: 10.1.11.0 cd25/1
Nexthop global: :: cd25/1
Nexthop local: :: cd25/1
BGP connection: non shared network

BGP neighbor is 10.1.12.1, remote AS 6500001, local AS 6500001, internal link, peer index: 4
  BGP version 4, local router ID 223.223.223.1, remote router ID 0.0.0.0
  BGP state = Idle, Socket state = Idle
  Last read Never, hold time is 90, keepalive interval is 30 seconds
  Received 0 messages, 0 notifications, 0 in queue
  Sent 0 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 0
  Index 2, Offset 0, Mask 0x4
  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 0; dropped 0
TCP MSS: (0), Advertise TCP MSS: (0), Send TCP MSS: (536),  Receive TCP MSS: (0)
Sock FD : (36)


BGP neighbor is 10.1.13.1, remote AS 6500001, local AS 6500001, internal link, peer index: 5
  BGP version 4, local router ID 223.223.223.1, remote router ID 0.0.0.0
  BGP state = Idle, Socket state = Idle
  Last read Never, hold time is 90, keepalive interval is 30 seconds
  Received 126 messages, 1 notifications, 0 in queue
  Sent 119 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 0
  Index 3, Offset 0, Mask 0x8
  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 1; dropped 1
TCP MSS: (0), Advertise TCP MSS: (0), Send TCP MSS: (536),  Receive TCP MSS: (0)
Sock FD : (37)
Last Reset: 04:49:44, due to BGP Notification received
Notification Error Message: (Cease/Peer Unconfigured.)


R1#
```

show bgp ipv6

```
R1#show bgp ipv6
BGP table version is 2, local router ID is 223.223.223.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
              l - labeled, S Stale, x-EVPN, w-Stitched-routes
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
Description : Ext-Color - Extended community color

    Network          Next Hop                    Metric    LocPrf     Weight Path
*>    10::/64          ::                             0        100       32768  ?        -
* i                    10::13(fe80::eac5:7aff:feb2:d1d)
                                                     0        100       0    ?        -

Total number of prefixes 1
R1#
```

show ipv6 route database bgp

```
R1#show ipv6 route database bgp
IP Route Table for VRF "default"
B   *>  10:10::/64 [200/0] via fe80::5e07:58ff:fe54:28a2, cd24/4, installed 00:01:27, last update
00:01:27 ago

Total number of IPv6 routes 1
R1#
```

show ipv6 route database

```
R1#show ipv6 route database
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
       > - selected route, * - FIB route,p - stale info
Timers: Uptime

IP Route Table for VRF "default"
C   *>  ::1/128 via ::, lo, installed 00:24:11, last update 00:24:11 ago
C   *>  10::/64 via ::, cd24/1, installed 00:03:19, last update 00:03:19 ago
C   *>  fe80::/64 via ::, cd24/1, installed 00:03:33, last update 00:03:33 ago
C       fe80::/64 via ::, cd25/3, installed 00:23:49, last update 00:23:49 ago
C       fe80::/64 via ::, cd24/4, installed 00:23:49, last update 00:23:49 ago
C       fe80::/64 via ::, cd25/1, installed 00:23:49, last update 00:23:49 ago

Total number of IPv6 routes 6
R1#
```

show ipv6 route

```
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,
 P - SRV6-POLICY,
        v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 00:23:49, last update 00:23:49 ago
C       10::/64 via ::, cd24/1, installed 00:02:57, last update 00:02:57 ago
C       fe80::/64 via ::, cd24/1, installed 00:03:11, last update 00:03:11 ago
R1#
```

VPNv4 NLRI with IPv6 Nexthop

Multi-protocol BGP (MP-BGP) specifies that the set of usable next-hop address families is determined by the Address Family Identifier (AFI) and the Subsequent Address Family Identifier (SAFI). Historically the AFI/SAFI definitions for the IPv4 address family only have provisions for advertising a Next Hop address that belongs to the IPv4 protocol when advertising IPv4 or VPN-IPv4 Network Layer Reachability Information (NLRI).

The extensions necessary to allow advertising IPv4 NLRI or VPN-IPv4 NLRI with a Next Hop address that belongs to the IPv6 protocol. This comprises an extension of the AFI/SAFI definitions to allow the address of the Next Hop for IPv4 NLRI or VPN-IPv4 NLRI to also belong to the IPv6 Protocol. The encoding of the Next Hop to determine which of the protocols the address actually belongs to, and a new BGP Capability allowing MP-BGP Peers to dynamically discover whether they can exchange IPv4 NLRI and VPN-IPv4 NLRI with an IPv6 Next Hop.

Topology

Figure 75. VPNv4 NLRI with IPv6 Nexthop

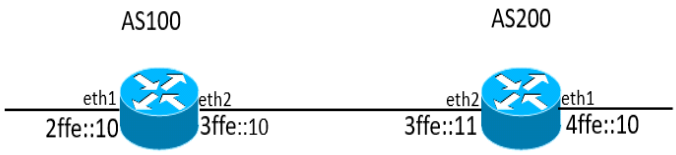

Configuration

R1

`#configure terminal`	Enter Configure mode.
`(config)# ip vrf VRF1`	Create a VRF, vrf1.
`(config-vrf)#rd 100:10`	Configure a route distinguisher value.
`(config-vrf)#route-target both 100:10`	Configure a route target both value to VRF.
`(config-vrf)#exit`	Exit from VRF configuration mode.
`(config)#interface eth1`	Enter Interface configuration mode.
`(config-if)#ip vrf forwarding VRF1`	Configure the interface to a VRF.
`(config-if)#ip address 30.30.30.30/24`	Assign an IP address to the interface.
`(config-if)#ipv6 address 2ffe::10/64`	Assign an IPV6 address to the interface.
`(config-if)#exit`	Exit from interface configuration mode.
`(config)#ip route vrf VRF1 50.50.50.0/24 eth1`	Create a VRF static route.
`(config)#interface eth2 Enter`	
`(config-if)#ipv6 address 3ffe::10/64`	Assign an IPV6 address to the interface.
`(config-if)#exit`	Exit from interface configuration mode.

© 2026 IP Infusion Inc. Proprietary

(config)#router bgp 100 Define the routing process.	
(config-router)#bgp router-id 1.1.1.1	Configure a fixed Router ID (1.1.1.1) for the BGP4+ routing process.
(config-router)#neighbor 2.2.2.2 remote-as 200	Define BGP neighbors, and establish a TCP session. 2.2.2.2 is the IP address of the neighbor (R1), and 200 is the neighbor's AS number.
(config-router)#neighbor 3ffe::11 remote-as 200	Define the BGP neighbor (R2), and establish a TCP session by specifying the IPv6 address (3ffe::11)and the AS number(200)of neighbor R2.
(config-router)#address-family vpnv4 unicast	Enter Address Family mode for configuring routing sessions that use IPv6 address prefixes.
(config-router-af)#neighbor 2.2.2.2 activate	Activate the neighbor R2 (2.2.2.2), and enable exchange of IP address prefix types with this neighbor.
(config-router-af)#neighbor 2.2.2.2 allow-ebgp-vpn	Configures the neighbor address to allow EBGP VPN.
(config-router-af)#neighbor 3ffe::11 allow-ebgp-vpn	Apply allow-ebgp-vpn to allow an eBGP neighbor to be a VPN peer.
(config-router-af)#neighbor 3ffe::11 activate	Activate the neighbor R2(3ffe::11), and enable exchange of IPv6 address prefix types with this neighbor.
(config-router-af)#neighbor 3ffe::11 capability extended-nexthop-encode	Enable capability to encode IPv6 nexthop in VPNv4 NLRI.
(config-router-af)#exit-address-family	Exit Address Family mode, and return to Router mode.
(config-router)#address-family ipv4 vrf VRF1	Enter IPv4 VRF address family.
(config-router-af)#redistribute static	Redistribute static into the IPv4 VRF address family.
(config-router-af)#exit-address-family	Exit Address Family mode, and return to Router mode.
(config-router)#exit	Exit Router mode, and return to Configure mode.
(config)#commit	Apply commit.
(config)#exit	Exit the config mode.

R2

#configure terminal	Enter Configure mode.
(config)# ip vrf VRF1	Create a VRF, vrf1.
(config-vrf)#rd 100:10	Configure a route distinguisher value.
(config-vrf)#route-target both 100:10	Configure a route target both value to VRF.
(config-vrf)#exit	Exit from VRF configuration mode.

© 2026 IP Infusion Inc. Proprietary

(config)#interface eth1	Enter Interface configuration mode.
(config-if)#ip vrf forwarding VRF1	Configure the interface to a VRF.
(config-if)#ip address 40.40.40.40/24	Assign an IP address to the interface.
(config-if)#ipv6 address 4ffe::10/64	Assign an IPV6 address to the interface.
(config-if)#exit	Exit from interface configuration mode.
(config)#ip route vrf VRF1 66.66.66.0/24 eth1	Create a VRF static route.
(config)#interface eth2	Enter Interface configuration mode.
(config-if)#ipv6 address 3ffe::11/64	Assign an IPV6 address to the interface.
(config-if)#exit	Exit from interface configuration mode.
(config)#router bgp 200 Define the routing process.	The number 200 specifies the AS number of R1.
(config-router)#bgp router-id 2.2.2.2	Configure a fixed Router ID (2.2.2.2) for the BGP4+ routing process.
(config-router)#neighbor 1.1.1.1 remote-as 100	Define BGP neighbors, and establish a TCP session. 1.1.1.1 is the IP address of the neighbor (R1), and 100 is the neighbor's AS number.
(config-router)#neighbor 3ffe::10 remote-as 100	Define the BGP neighbor (R2), and establish a TCP session by specifying the IPv6 address (3ffe::10)and the AS number(100)of neighbor R2.
(config-router)#address-family vpnv4 unicast	Enter Address Family mode for configuring routing sessions that use IPv6 address prefixes.
(config-router-af)#neighbor 1.1.1.1 activate	Activate the neighbor R2(1.1.1.1), and enable exchange of IP address prefix types with this neighbor.
(config-router-af)#neighbor 1.1.1.1 allow-ebgp-vpn	Configures the neighbor address to allow EBGP VPN.
(config-router-af)#neighbor 3ffe::10 allow- ebgp-vpn	Apply allow-ebgp-vpn to allow an eBGP neighbor to be a VPN peer.
(config-router-af)#neighbor 3ffe::10 activate	Activate the neighbor R2(3ffe::10), and enable exchange of IPv6 address prefix types with this neighbor.
(config-router-af)#neighbor 3ffe::10 capability extended-nexthop-encode	Enable capability to encode IPv6 nexthop in VPNv4 NLRI.
(config-router-af)#exit-address-family	Exit Address Family mode, and return to Router mode.
(config-router)#address-family ipv4 vrf VRF1	Enter IPv4 VRF address family.
(config-router-af)#redistribute static	Redistribute static into the IPv4 VRF address family.
(config-router-af)#exit-address-family	Exit Address Family mode, and return to Router mode.
(config-router)#exit	Exit Router mode, and return to Configure mode.

© 2026 IP Infusion Inc. Proprietary

(config)#commit	Apply commit.
(config)#exit	Exit the config mode.

Validation

R1

```
#show running-config router bgp
router bgp 100
 bgp router-id 1.1.1.1
 neighbor 2.2.2.2 remote-as 200
 neighbor 3ffe::11 remote-as 200
!
 address-family vpnv4 unicast
 neighbor 2.2.2.2 activate
 neighbor 3ffe::11 allow-ebgp-vpn
 neighbor 3ffe::11 activate
 neighbor 3ffe::11 capability extended-nexthop-encode
 exit-address-family
!
 address-family ipv4 vrf vrf1
 redistribute static
 exit-address-family
 #sh ip bgp neighbors
BGP neighbor is 3ffe::11, remote AS 200, 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 00:04:59
  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 VPNv4 Unicast: advertised and received
    Address family IPv4 Unicast: advertised and received
  Received 167 messages, 1 notifications, 0 in queue
  Sent 166 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: VPNv4 Unicast
  BGP table version 1, neighbor version 1
  Index 3, Offset 0, Mask 0x8
    Extended Nexthop Encoding: advertised and received
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes
 Connections established 1; dropped 0
Local host: 3ffe::10, Local port: 60148
Foreign host: 3ffe::11, Foreign port: 179
Nexthop: 2.2.2.2
Nexthop global: 3ffe::10
Nexthop local: fe80::ba6a:97ff:fee0:61be
BGP connection: shared network
Last Reset: 00:05:04, due to Configuration Change (Cease Notification sent)
Notification Error Message: (Cease/Other Configuration Change.)

#sh ip bgp vpnv4 all
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
Route Distinguisher: 100:10 (Default for VRF vrf1)
*>   55.55.55.0/24    0.0.0.0              0        100       32768  ?
*>   66.66.66.0/24    0.0.0.0              0        100       0      ?
```

```
 Announced routes count = 2
 Accepted routes count = 0
Route Distinguisher: 100:10
* i  55.55.55.0/24    3ffe::11              0        100       0    ?
* i  66.66.66.0/24    3ffe::11              0        100       0    ?
 Announced routes count = 0
 Accepted routes count = 2

#sh ip bgp vpnv4 all summary
BGP router identifier 2.2.2.2, 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/Dow
n   State/PfxRcd
2.2.2.2                  4   400    0          0       0      0      0     never
        Idle

3ffe::11                 4   400    163       163       1      0      0  00:00:04
             2
Total number of neighbors 2

Total number of Established sessions 1
```

R2

```
#show running-config router bgp
router bgp 200
 bgp router-id 2.2.2.2
 neighbor 1.1.1.1 remote-as 100
 neighbor 3ffe::10 remote-as 100
!
 address-family vpnv4 unicast
 neighbor 1.1.1.1 activate
 neighbor 3ffe::10 allow-ebgp-vpn
 neighbor 3ffe::10 activate
 neighbor 3ffe::10 capability extended-nexthop-encode
 exit-address-family
!
 address-family ipv4 vrf vrf1
 redistribute static
 exit-address-family
 #sh ip bgp neighbors
BGP neighbor is 3ffe::10, remote AS 100, local AS 200, external link
  BGP version 4, local router ID 2.2.2.2, remote router ID 1.1.1.1
  BGP state = Established, up for 00:04:59
  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 VPNv4 Unicast: advertised and received
    Address family IPv4 Unicast: advertised and received
  Received 167 messages, 1 notifications, 0 in queue
  Sent 166 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: VPNv4 Unicast
  BGP table version 1, neighbor version 1
  Index 3, Offset 0, Mask 0x8
    Extended Nexthop Encoding: advertised and received
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes
 Connections established 1; dropped 0
```

```
Local host: 3ffe::11, Local port: 60148
Foreign host: 3ffe::10, Foreign port: 179
Nexthop: 1.1.1.1
Nexthop global: 3ffe::11
Nexthop local: fe80::ba6a:97ff:fee0:61ab
BGP connection: shared network
Last Reset: 00:05:04, due to Configuration Change (Cease Notification sent)
Notification Error Message: (Cease/Other Configuration Change.)
#sh ip bgp vpnv4 all
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
Route Distinguisher: 100:10 (Default for VRF vrf1)
*>   55.55.55.0/24    0.0.0.0              0        100       32768  ?
*>   66.66.66.0/24    0.0.0.0              0        100       0      ?
 Announced routes count = 2
 Accepted routes count = 0
Route Distinguisher: 100:10
* i  55.55.55.0/24    3ffe::10              0        100       0    ?
* i  66.66.66.0/24    3ffe::10              0        100       0    ?
 Announced routes count = 0
 Accepted routes count = 2

#sh ip bgp vpnv4 all summary
BGP router identifier 1.1.1.1, 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/Dow
n   State/PfxRcd
1.1.1.1                  4   400    0          0       0      0      0     never
        Idle

3ffe::10                 4   400    163       163       1      0      0  00:00:04
              2
Total number of neighbors 2

Total number of Established sessions 1
```

BGP Labeled Unicast

As well as distributing routes, BGP with Multiprotocol Extensions (MP-BGP) can advertise MPLS label mappings that are mapped to routes. BGP Labeled Unicast (BGP-LU) attaches an MPLS label to an advertised IGP prefix and distributes the MPLS label mapped to the prefix to its peers.

With BGP-LU, a network can be divided into multiple regions to limit the total number of LSPs and enable failures to be contained and restored in a single region These regions operate separate instances of the IGP and use BGP-LU to advertise route information between inter-region routers.

A configuration for BGP-LU uses these type of nodes:

- Provider Edge (PE) nodes advertise label bindings to remote PEs in other regions. These advertisements only affect the PE routers and the ABRs and not provider routers (“P”) in the core network.
- Area Border Router (ABR) nodes advertise the label bindings to remote PEs in other regions.

Notes:

- Default route advertisement is not supported in BGP Labeled Unicast. Because BGP-LU operates as both a service and transport layer, default routes at the LU level in MPLS can create ambiguity in route resolution and dependency tracking. OcNOS prevents such configurations to maintain routing consistency. To avoid this, do not advertise the default route in LU, or apply a route map to filter it (deny default route using route map) on LU neighbors.
- The default behavior in BGP LU for all connected or local IPv4 unicast routes advertised in labeled-unicast address-family mode is modified to assign implicit null label. For more information, refer to BGP Labeled Unicast with Implicit Null Label for Local Routes (page 720).

Hierarchical BGP LU: BGP Labeled Unicast (LU) deployment is not supported as both transport or underlay and service or overlay within a single OcNOS node. This refers to configuring an LU LSP to utilize another LU LSP as its service transport. The MPLS service manager in OcNOS restricts LU FTN over LU FTN. However, hierarchical Labeled BGP is supported, allowing for L3VPN, 6PE, or 6VPE LSPs over LU LSPs. LU LSPs can function as a service or transport and are compatible with various underlay labeling protocols.

© 2026 IP Infusion Inc. Proprietary

BGP Labeled Unicast as Transport Configuration

Topology

Figure 76. BGP labeled unicast

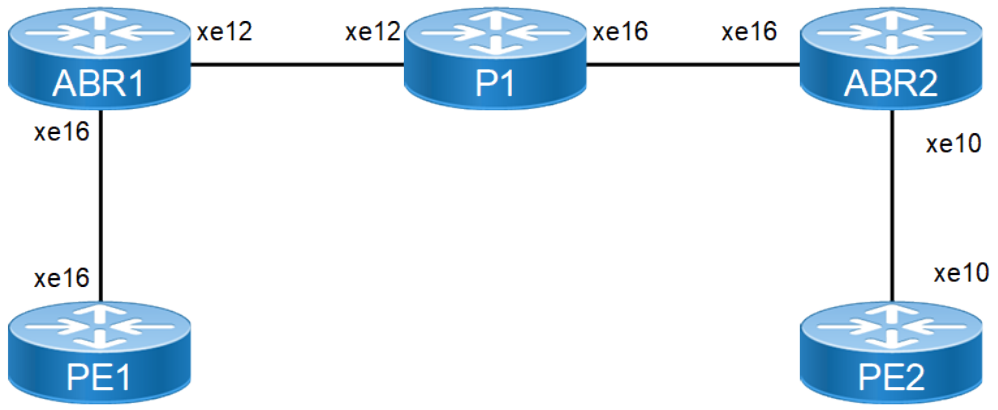

PE1 Configuration

The following are the PE1 configuration procedures.

`#configure terminal`	Enter the Configure mode.
`(config)#interface lo`	Enter interface mode
`(config-if)#ip address 11.11.11.55/32 secondary`	Configure the IP address of the interface loopback
`(config-if)#exit`	Exit interface mode
`(config)#interface xe16`	Enter interface mode
`(config-if)#ip address 172.4.5.55/24`	Configure the IP address of the interface eth1
`(config-if)#label-switching`	Enable label-switching on interface eth1
`(config-if)#exit`	Exit interface mode.
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router ospf 1`	Configure the routing process OSPF with process Id 1.
`(config-router)#network 172.4.5.0/24 area 0`	Define the interface (172.4.5.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#network 11.11.11.55/32 area 0`	Define the interface (11.11.11.55/32) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).

© 2026 IP Infusion Inc. Proprietary

`(config-router)#exit`	Exit from router ospf mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router bgp 100`	Enter Router BGP mode
`(config-router)#neighbor 21.21.21.56 update-source lo`	Add loopback IP of PE2 as neighbor with neighbor AS
`(config-router)#neighbor 21.21.21.56 update-source lo`	Update the source for that particular neighbor as loopback interface
`(config-router)#neighbor 172.4.5.52 remote-as 100`	Add neighbor with neighbor AS
`(config-router)#allocate-label all`	Allocate labels
`(config-router)#address-family ipv4 labeled-unicast`	Enter into labeled-unicast address family
`(config-router-af)#neighbor 172.4.5.52 activate`	Activate the neighbor inside labeled-unicast address family
`(config-router-af)#exit-address-family`	Exit from address family IPv4 labeled unicast
`(config-router)#address-family vpnv4 unicast`	Enter into vpnv4 unicast address family
`(config-router-af)#neighbor 21.21.21.56 activate`	Activate the neighbor inside vpnv4 address family
`(config-router-af)#exit-address-family`	Exit from address family vpnv4.
`(config-router)#address-family ipv4 unicast`	Enter into ipv4 unicast address family
`(config-router-af)#network 11.11.11.55/32`	Advertise the loopback of RTR1 in BGP
`(config-router-af)#exit-address-family`	Exit from address family vpnv4
`(config-router)#exit`	Exit from router BGP mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)# ip vrf vrf1`	Specify the name of the VRF (vrf1) 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 100:300`	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 xe10`	Enter interface mode
`(config-if)# ip vrf forwarding vrf1`	Bind the interface (eth2) to the VRF vrf1
`(config-if)# ip address 172.10.20.55/24`	Configure the IP address of the interface eth2
`(config-if)#exit`	Exit interface mode.
`(config)#router bgp 100`	Enter Router BGP mode

© 2026 IP Infusion Inc. Proprietary

(config-router)# address-family ipv4 vrf vrf1	Enter address family ipv4 vrf mode
(config-router-af)# redistribute connected	Redistribute connected routes
(config-router-af)# exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

ABR1 Configuration

The following are the ABR1 configuration procedures.

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 22.22.22.52/32 secondary	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)#interface xe16	Enter interface mode
(config-if)#ip address 172.4.5.52/24	Configure the IP address of the interface eth1
(config-if)#label-switching	Enable label-switching on interface eth1
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#interface xe12	Enter interface mode
(config-if)#ip address 172.6.7.52/24	Configure the IP address of the interface eth1
(config-if)#label-switching	Enable label-switching on interface eth1
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router ospf 1	Configure the routing process OSPF with process Id 1
(config-router)#network 172.4.5.0/24 area 0	Define the interface (172.4.5.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
(config-router)#network 172.6.7.0/24 area 0	Define the interface (172.6.7.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
(config-router)#exit	Exit from router ospf mode
(config)#router bgp 100	Enter Router BGP mode

© 2026 IP Infusion Inc. Proprietary

(config-router)# neighbor 172.4.5.55 remote-as 100	Add neighbor with neighbor AS
(config-router)# neighbor 172.6.7.54 remote-as 100	Add neighbor with neighbor AS
(config-router)#allocate-label all	Allocate labels
(config-router)# address-family ipv4 labeled-unicast	Enter into labeled-unicast address family
(config-router-af)#neighbor 172.6.7.54 activate	Activate the neighbor inside labeled-unicast address family
(config-router-af)#neighbor 172.4.5.55 activate	Activate the neighbor inside labeled-unicast address family
(config-router-af)# neighbor 172.4.5.55 route-reflector-client	Enable Route reflector client for the neighbor inside address family IPv4 labeled unicast
(config-router-af)# neighbor 172.6.7.54 route-reflector-client	Enable Route reflector client for the neighbor inside address family IPv4 labeled unicast
(config-router-af)# neighbor 172.4.5.55 next-hop-self	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast
(config-router-af)# neighbor 172.6.7.54 next-hop-self	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast
(config-router-af)#exit-address-family	Exit address-family mode
(config-router)#exit	Exit from router BGP mode
(config)#commit	Commit the candidate configuration to the running configuration.

P1 Configuration

The following are the P1 configuration procedures.

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)# ip address 200.200.200.54/32 secondary	Configure the IP address of the interface loopbak
(config-if)#exit	Exit interface mode
(config)#interface xe16	Enter interface mode
(config-if)#ip address 172.1.2.54/24	Configure the IP address of the interface eth1
(config-if)#label-switching	Enable label-switching on interface eth1
(config-if)#exit	Exit interface mode.
(config)#interface xe12	Enter interface mode
(config-if)#ip address 172.6.7.54/24	Configure the IP address of the interface eth1
(config-if)#label-switching	Enable label-switching on interface eth1
(config-if)#exit	Exit interface mode.

© 2026 IP Infusion Inc. Proprietary

`(config)#router ospf 1`	Configure the routing process OSPF with process Id 1
`(config-router)#network 172.1.2.0/24 area 0`	Define the interface (172.1.2.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#network 172.6.7.0/24 area 0`	Define the interface (172.6.7.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#exit`	Exit from router ospf mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router bgp 100`	Enter Router BGP mode
`(config-router)# neighbor 172.1.2.53 remote-as 100`	Add neighbor with neighbor AS
`(config-router)# neighbor 172.6.7.52 remote-as 100`	Add neighbor with neighbor AS
`(config-router)#allocate-label all`	Allocate labels
`(config-router)# address-family ipv4 labeled-unicast`	Enter into labeled-unicast address family
`(config-router-af)#neighbor 172.6.7.52 activate`	Activate the neighbor inside labeled-unicast address family
`(config-router-af)#neighbor 172.1.2.53 activate`	Activate the neighbor inside labeled-unicast address family
`(config-router-af)# neighbor 172.1.2.53 route-reflector-client`	Enable Route reflector client for the neighbor inside address family IPv4 labeled unicast
`(config-router-af)# neighbor 172.6.7.52 route-reflector-client`	Enable Route reflector client for the neighbor inside address family IPv4 labeled unicast
`(config-router-af)# neighbor 172.1.2.53 next-hop-self`	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast
`(config-router-af)# neighbor 172.6.7.52 next-hop-self`	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast
`(config-router)# exit-address-family`	Exit from address family.
`(config-router)#exit`	Exit from router BGP mode
`(config)#commit`	Commit the candidate configuration to the running configuration.

ABR2 Configuration

The following are the ABR2 configuration procedures.

`#configure terminal`	Enter the Configure mode.

`(config)#interface lo`	Enter interface mode
`(config-if)# ip address 44.44.44.53/32 secondary`	Configure the IP address of the interface loopback
`(config-if)#exit`	Exit interface mode
`(config)#interface xe16`	Enter interface mode
`(config-if)#ip address 172.1.2.53/24`	Configure the IP address of the interface eth1
`(config-if)#label-switching`	Enable label-switching on interface eth1
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe10`	Enter interface mode
`(config-if)#ip address 172.3.4.53/24`	Configure the IP address of the interface eth1
`(config-if)#label-switching`	Enable label-switching on interface eth1
`(config-if)#exit`	Exit interface mode.
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router ospf 1`	Configure the routing process OSPF with process Id 1
`(config-router)#network 172.1.2.0/24 area 0`	Define the interface (172.1.2.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#network 172.3.4.0/24 area 0`	Define the interface (172.3.4.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#exit`	Exit from router ospf mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router bgp 100`	Enter Router BGP mode
`(config-router)# neighbor 172.1.2.54 remote-as 100`	Add neighbor with neighbor AS
`(config-router)# neighbor 172.3.4.56 remote-as 100`	Add neighbor with neighbor AS
`(config-router)#allocate-label all`	Allocate labels
`(config-router)# address-family ipv4 labeled-unicast`	Enter into labeled-unicast address family
`(config-router-af)#neighbor 172.3.4.56 activate`	Activate the neighbor inside labeled-unicast address family
`(config-router-af)#neighbor 172.1.2.54 activate`	Activate the neighbor inside labeled-unicast address family
`(config-router-af)# neighbor 172.3.4.56 route-reflector-client`	Enable Route reflector client for the neighbor inside address family IPv4 labeled unicast
`(config-router-af)# neighbor 172.1.2.54 route-`	Enable Route reflector client for the neighbor

© 2026 IP Infusion Inc. Proprietary

reflector-client	inside address family IPv4 labeled unicast
(config-router-af)# neighbor 172.3.4.56 next-hop-self	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast
(config-router-af)# neighbor 172.1.2.54 next-hop-self	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast
(config-router-af)#exit-address-family	Exit from address family vpnv4
(config-router)#exit	Exit from router BGP mode
(config)#end	Exit from config mode
(config)#commit	Commit the candidate configuration to the running configuration.

PE2 Configuration

The following are the PE2 configuration procedures.

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 21.21.21.56/32 secondary	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)#interface xe10	Enter interface mode
(config-if)#ip address 172.3.4.56/24	Configure the IP address of the interface eth1
(config-if)#label-switching	Enable label-switching on interface eth1
(config-if)#exit	Exit interface mode.
(config)#router ospf 1	Configure the routing process OSPF with process Id 1
(config-router)#network 172.3.4.0/24 area 0	Define the interface (172.3.4.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
(config-router)#network 21.21.21.56/32 area 0	Define the interface (21.21.21.56/32) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
(config-router)#exit	Exit from router ospf mode
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router bgp 100	Enter Router BGP mode
(config-router)#neighbor 11.11.11.55 remote-as 100	Add loopback ip of PE1 as neighbor with neighbor AS

© 2026 IP Infusion Inc. Proprietary

`(config-router)#neighbor 11.11.11.55 update-source lo`	Update the source for that particular neighbor as loopback interface
`(config-router)# neighbor 172.3.4.53 remote-as 100`	Add neighbor with neighbor AS
`(config-router)#allocate-label all`	Allocate labels
`(config-router)# address-family ipv4 labeled-unicast`	Enter into labeled-unicast address family
`(config-router-af)#neighbor 172.3.4.53 activate`	Activate the neighbor inside labeled-unicast address family
`(config-router-af)#exit-address-family`	Exit from address family IPv4 labeled unicast
`(config-router)#address-family vpnv4 unicast`	Enter into vpnv4 unicast address family
`(config-router-af)#neighbor 11.11.11.55 activate`	Activate the neighbor inside vpnv4 address family
`(config-router-af)#exit-address-family`	Exit from address family vpnv4
`(config-router-af)#network 21.21.21.56/32`	Advertise the loopback in BGP
`(config-router-af)#exit-address-family`	Exit from address family ipv4 unicast
`(config-router)#exit`	Exit from router BGP mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)# ip vrf vrf1`	Specify the name of the VRF (vrf1) 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 100:300`	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 xe0`	Enter interface mode
`(config-if)# ip vrf forwarding vrf1`	Bind the interface (eth2) to the VRF vrf1
`(config-if)# ip address 172.23.4.56/24`	Configure the IP address of the interface eth2
`(config-if)#exit`	Exit interface mode.
`(config)#router bgp 100`	Enter Router BGP mode
`(config-router)# address-family ipv4 vrf vrf1`	Enter address family ipv4 vrf mode
`(config-router)# redistribute connected`	Redistribute connected routes
`(config-router)#exit`	Exit from router mode into privilege mode
`(config)#commit`	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

BGP LU Transport Validation

PE1

```
PE1#show ip bgp neighbors 21.21.21.56
BGP neighbor is 21.21.21.56, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 11.11.11.55, remote router ID 21.21.21.56
  BGP state = Established, up for 00:01:32
  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 VPNv4 Unicast: advertised and received
  Received 7 messages, 0 notifications, 0 in queue
  Sent 9 messages, 0 notifications, 0 in queue
  Route refresh request: received 1, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo
 For address family: VPNv4 Unicast
  BGP table version 2, neighbor version 2
  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: 11.11.11.55, Local port: 179
Foreign host: 21.21.21.56, Foreign port: 38231
Nexthop: 11.11.11.55
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network


PE1#show ip bgp neighbors 172.4.5.52
BGP neighbor is 172.4.5.52, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 11.11.11.55, remote router ID 22.22.22.52
  BGP state = Established, up for 00:26:00
  Last read 00:00:07, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: received
    Address family IPv4 Labeled-Unicast: advertised and received
  Received 67 messages, 2 notifications, 0 in queue
  Sent 70 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 Labeled-Unicast
  BGP table version 2, neighbor version 2
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  1 announced prefixes

 Connections established 3; dropped 2
Local host: 172.4.5.55, Local port: 179
Foreign host: 172.4.5.52, Foreign port: 34324
Nexthop: 172.4.5.55
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:26:00, due to Administratively Reset (Cease Notification sent)
Notification Error Message: (Cease/Administratively Reset.)

PE1#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
S - stale
```

```
     Network          Next Hop        In Label     Out Label
*>  11.11.11.55/32   0.0.0.0           24320          -
*>i 21.21.21.56/32   172.4.5.52        24322        24321

PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
       B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

Code    FEC                 FTN-ID    Nhlfe-ID  Tunnel-id   Pri   LSP-Type     O
ut-Label    Out-Intf    ELC   Nexthop
   B>   21.21.21.56/32      1         2         -           -     LSP_DEFAULT  2
4321        xe16        No    172.4.5.52
PE1#

PE1#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
        K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP , K - CLI , V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

Code    FEC/VRF/L2CKT    ILM-ID      In-Label    Out-Label   In-Intf    Out-Intf
/VRF       Nexthop                  LSP-Type
   B>   vrf1               2          24321       Nolabel     N/A        vrf1
           N/A                      LSP_DEFAULT
   B>   11.11.11.55/32     1          24320       Nolabel     N/A        N/A
           127.0.0.1                LSP_DEFAULT
   B>   21.21.21.56/32     3          24322       24321       N/A        N/A
           172.4.5.52               LSP_DEFAULT
PE1#
```

ABR1

```
ABR1#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface           Instance ID
11.11.11.55       1   Full/DR         00:00:35    172.4.5.55      xe16                    0
200.200.200.54    1   Full/Backup     00:00:32    172.6.7.54      xe12                    0
ABR1#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 11.11.11.55/32   172.4.5.55        24320        24320
*>i 21.21.21.56/32   172.6.7.54        24321        24321

ABR1#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 11.11.11.55/32   172.4.5.55        24320        24320
*>i 21.21.21.56/32   172.6.7.54        24321        24321

ABR1#show mpls forwarding-table | include 11.11.11.55
   B>    11.11.11.55/32      1        0           Yes   LSP_
DEFAULT     24320        xe16        No        172.4.5.55
```

P1

```
P1#show ip bgp neighbors 172.6.7.52
BGP neighbor is 172.6.7.52, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 200.200.200.54, remote router ID 22.22.22.52
  BGP state = Established, up for 01:33:27
  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: received
    Address family IPv4 Labeled-Unicast: advertised and received
  Received 229 messages, 0 notifications, 0 in queue
  Sent 227 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 Labeled-Unicast
  BGP table version 7, neighbor version 7
  Index 1, Offset 0, Mask 0x2
  Route-Reflector Client
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  1 announced prefixes

 Connections established 3; dropped 2
Local host: 172.6.7.54, Local port: 47560
Foreign host: 172.6.7.52, Foreign port: 179
Nexthop: 172.6.7.54
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 01:33:32, due to Configuration Change (Cease Notification sent)
Notification Error Message: (Cease/Other Configuration Change.)

P1#show ip bgp neighbors 172.1.2.53
BGP neighbor is 172.1.2.53, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 200.200.200.54, remote router ID 44.44.44.53
  BGP state = Established, up for 01:35:08
  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 Labeled-Unicast: advertised and received
  Received 230 messages, 2 notifications, 0 in queue
  Sent 233 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 Labeled-Unicast
  BGP table version 7, neighbor version 7
  Index 2, Offset 0, Mask 0x4
  Route-Reflector Client
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  1 announced prefixes

 Connections established 3; dropped 2
Local host: 172.1.2.54, Local port: 179
Foreign host: 172.1.2.53, Foreign port: 58824
Nexthop: 172.1.2.54
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 01:35:13, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)

P1#show ip bgp labeled-unicast
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 11.11.11.55/32   172.6.7.52        24320        24320
*>i 21.21.21.56/32   172.1.2.53        24321        24321

P1#show mpls forwarding-table | include 11.11.11.55
   B>   11.11.11.55/32      1         1         -           -     LSP_
DEFAULT  24320        xe12        No    172.6.7.52
```

ABR2

```
ABR2#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
S - stale
     Network          Next Hop        In Label     Out Label
*>i 11.11.11.55/32   172.1.2.54        24320        24320
*>i 21.21.21.56/32   172.3.4.56        24321        24321

ABR2#show mpls forwarding-table | include 11.11.11.55
   B>   11.11.11.55/32      1         1         -           -     LSP_
DEFAULT  24320        xe16        No    172.1.2.54
```

PE2

```
PE2#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 11.11.11.55/32   172.3.4.53        24320        24320
*>  21.21.21.56/32   0.0.0.0           24321          -

PE2#show mpls forwarding-table | include 11.11.11.55
   B>   11.11.11.55/32      1         1         -           -     LSP_
DEFAULT  24320        xe10        No    172.3.4.53

PE2#show mpls ftn-table
 Primary FTN entry with FEC: 11.11.11.55/32, id: 1, row status: Active, Tunnel-Policy: N/A
  Owner: BGP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
     Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: BGP, Stale: NO, out intf: xe10, out label: 24320
    Nexthop addr: 172.3.4.53        cross connect ix: 1, op code: Push

PE2#show mpls vrf-table
Output for IPv4 VRF table with id: 2
 Primary FTN entry with FEC: 172.10.20.0/24, id: 1, row status: Active, Tunnel-Policy: N/A
  Owner: BGP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Transport Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
     Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 2
      Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 2, owner: BGP, Stale: NO, BGP out intf: xe10, transport out intf: xe10,
out label: 24321
    Nexthop addr: 11.11.11.55        cross connect ix: 3, op code: Push and Lookup


PE2#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
        K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
```

```
       B - BGP , K - CLI , V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

Code     FEC/VRF/L2CKT     ILM-ID      In-Label    Out-Label    In-Intf    Out-
Intf/VRF       Nexthop                    LSP-Type
   B>    21.21.21.56/32      2           24321       Nolabel      N/A        N/A              127.0.0.1
                LSP_DEFAULT
   B>    11.11.11.55/32      1           24320       24320        N/A        N/A              172.3.4.53
                LSP_DEFAULT
   B>    vrf1                3           24322       Nolabel      N/A        vrf1             N/A
                LSP_DEFAULT
```

BGP Labeled Unicast with Seamless MPLS Configuration

Topology

Figure 77. BGP_LU with MPLS

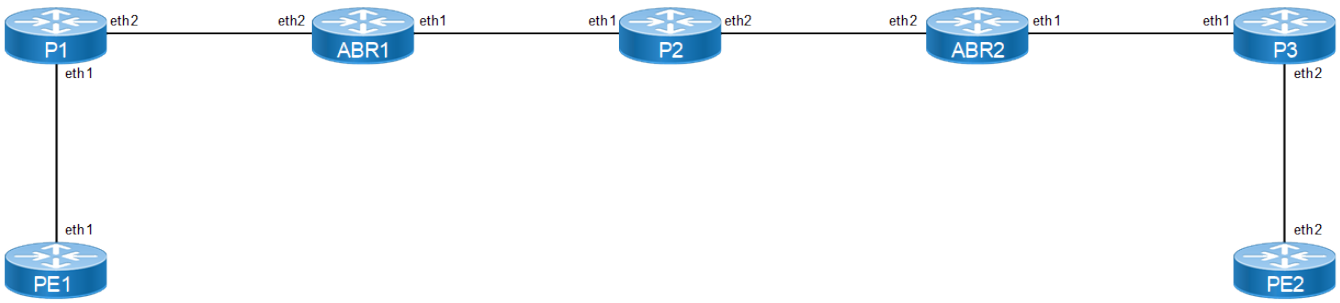

PE1 Configuration

The following are the PE1 configuration procedures.

`#configure terminal`	Enter the Configure mode.
`(config)#interface lo`	Enter interface mode
`(config-if)#ip address 1.1.1.54/32 secondary`	Configure the IP address of the interface loopback
`(config-if)#exit`	Exit interface mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip address 10.1.1.54/24`	Configure the IP address of the interface eth1
`(config-if)#label-switching`	Enable label-switching on interface eth1
`(config-if)# enable-ldp ipv4`	Enable LDP process on eth1 interface
`(config-if)#exit`	Exit interface mode.
`(config)#router ospf 10`	Configure the routing process OSPF with process id 10
`(config-router)#network 1.1.1.54/32 area 0`	Define the interface (1.1.1.54/32) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone

© 2026 IP Infusion Inc. Proprietary

	area).
(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)#exit	Exit from router ospf mode
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router bgp 100	Enter Router BGP mode
(config-router)#neighbor 7.7.7.55 remote-as 100	Add loopback ip of PE2 as neighbor with neighbor AS
(config-router)#neighbor 7.7.7.55 update-source lo	Update the source for that particular neighbor as loopback interface
(config-router)#neighbor 3.3.3.52 remote-as 100	Add loopback ip of ABR1 as neighbor with neighbor AS
(config-router)#neighbor 3.3.3.52 update-source lo	Update the source for that particular neighbor as loopback interface
(config-router)#allocate-label all	Allocate labels
(config-router)#address-family ipv4 labeled-unicast	Enter into labeled-unicast address family
(config-router-af)#neighbor 3.3.3.52 activate	Activate the neighbor inside labeled-unicast address family
(config-router-af)#exit-address-family	Exit from address family IPv4 labeled unicast
(config-router)#address-family vpnv4 unicast	Enter into vpnv4 unicast address family
(config-router-af)#neighbor 7.7.7.55 activate	Activate the neighbor inside vpnv4 address family
(config-router)#address-family ipv4 labeled-unicast	Enter into labeled-unicast address family
(config-router)#address-family ipv4 unicast	Enter into ipv4 unicast address family
(config-router-af)#network 1.1.1.54/32	Advertise the loopback of RTR1 in BGP
(config-router-af)#exit-address-family	Exit from address family
(config-router)#exit	Exit from router BGP mode
(config)# router ldp	Configure Router LDP instance
(config-router)# transport-address ipv4 1.1.1.54 0	Configure Transport address for LDP with label space value 0
(config-router)#exit	Exit from router mode
(config)#commit	Commit the candidate configuration to the running configuration.
(config)# ip vrf vrf1	Specify the name of the VRF (vrf1) 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 100:300	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 eth2	Enter interface mode
(config-if)# ip vrf forwarding vrf1	Bind the interface (eth2) to the VRF vrf1
(config-if)# ip address 61.1.1.54/24	Configure the IP address of the interface eth2
(config-if)#exit	Exit interface mode.
(config)#router bgp 100	Enter Router BGP mode
(config-router)# address-family ipv4 vrf vrf1	Enter address family ipv4 vrf mode
(config-router-af)# redistribute connected	Redistribute connected routes
(config-router-af)# exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

P1 Configuration

The following are the P1 configuration procedures.

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 2.2.2.23/32	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)#interface eth1	Enter interface mode
(config-if)#ip address 10.1.1.23/24	Configure the IP address of the interface eth1
(config-if)#label-switching	Enable label-switching on interface eth1
(config-if)# enable-ldp ipv4	Enable LDP process on eth1 interface
(config-if)#exit	Exit interface mode.
(config)#interface eth2	Enter interface mode
(config-if)#ip address 20.1.1.23/24	Configure the IP address of the interface eth2
(config-if)#label-switching	Enable label-switching on interface eth2
(config-if)# enable-ldp ipv4	Enable LDP process on eth2 interface
(config-if)#exit	Exit interface mode.
(config)#router ospf 10	Configure the routing process OSPF with process id 10

© 2026 IP Infusion Inc. Proprietary

`(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 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 2.2.2.23/32 area 0`	Define the interface (2.2.2.23/32) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#exit`	Exit from router ospf mode
`(config)# router ldp`	Configure Router LDP instance
`(config-router)# transport-address ipv4 2.2.2.23 0`	Configure Transport address for LDP with label space value 0
`(config-router)#exit`	Exit from router mode
`(config)# ip route 7.7.7.55/32 20.1.1.52`	Specify the destination prefix and mask for the network and a gateway.
`(config)#commit`	Commit the candidate configuration to the running configuration.

ABR1 Configuration

The following are the ABR1 configuration procedures.

`#configure terminal`	Enter the Configure mode.
`(config)#interface lo`	Enter interface mode
`(config-if)#ip address 3.3.3.52/32`	Configure the IP address of the interface loopback
`(config-if)#exit`	Exit interface mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip address 30.1.1.52/24`	Configure the IP address of the interface eth1
`(config-if)#label-switching`	Enable label-switching on interface eth1
`(config-if)# enable-ldp ipv4`	Enable LDP process on eth1 interface
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode
`(config-if)#ip address 20.1.1.52/24`	Configure the IP address of the interface eth2
`(config-if)#label-switching`	Enable label-switching on interface eth2
`(config-if)# enable-ldp ipv4`	Enable LDP process on eth2 interface

© 2026 IP Infusion Inc. Proprietary

`(config-if)#exit`	Exit interface mode.
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router ospf 20`	Configure the routing process OSPF with process id 20
`(config-router)#network 3.3.3.52/32 area 0`	Define the interface (3.3.3.52/32) 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).
`(config-router)#exit`	Exit from router ospf mode
`(config)# router ldp`	Configure Router LDP instance
`(config-router)# transport-address ipv4 3.3.3.52 0`	Configure Transport address for LDP with label space value 0
`(config-router)#exit`	Exit from router mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router bgp 100`	Enter Router BGP mode
`(config-router)#neighbor 1.1.1.54 remote-as 100`	Add loopback ip of PE2 as neighbor with neighbor AS
`(config-router)#neighbor 1.1.1.54 update-source lo`	Update the source for that particular neighbor as loopback interface
`(config-router)#neighbor 5.5.5.56 remote-as 100`	Add loopback ip of ABR1 as neighbor with neighbor AS
`(config-router)#neighbor 5.5.5.56 update-source lo`	Update the source for that particular neighbor as loopback interface
`(config-router)#allocate-label all`	Allocate labels
`(config-router)#address-family ipv4 labeled-unicast`	Enter into labeled-unicast address family
`(config-router-af)#neighbor 1.1.1.54 activate`	Activate the neighbor inside labeled-unicast address family
`(config-router-af)#neighbor 5.5.5.56 activate`	Activate the neighbor inside labeled-unicast address family
`(config-router-af)# neighbor 1.1.1.54 route-reflector-client`	Enable Route reflector client for the neighbor inside address family IPv4 labeled unicast
`(config-router-af)# neighbor 5.5.5.56 route-reflector-client`	Enable Route reflector client for the neighbor inside address family IPv4 labeled unicast
`(config-router-af)# neighbor 1.1.1.54 next-hop-self`	Enable next hop self for the particular neighbor

© 2026 IP Infusion Inc. Proprietary

	inside address family IPv4 labeled unicast
(config-router-af)# neighbor 5.5.5.56 next-hop-self	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast
(config-router-af)#exit-address-family	Exit from address family labeled-unicast
(config)#commit	Commit the candidate configuration to the running configuration.

P2 Configuration

The following are the P2 configuration procedures.

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 4.4.4.53/32	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)#interface eth1	Enter interface mode
(config-if)#ip address 30.1.1.53/24	Configure the IP address of the interface eth1
(config-if)#label-switching	Enable label-switching on interface eth1
(config-if)# enable-ldp ipv4	Enable LDP process on eth1 interface
(config-if)#exit	Exit interface mode.
(config)#interface eth2	Enter interface mode
(config-if)#ip address 40.1.1.53/24	Configure the IP address of the interface eth2
(config-if)#label-switching	Enable label-switching on interface eth2
(config-if)# enable-ldp ipv4	Enable LDP process on eth2 interface
(config-if)#exit	Exit interface mode
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router ospf 20	Configure the routing process OSPF with process id 20
(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 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)#network 4.4.4.53/32 area 0	Define the interface (4.4.4.53/32) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone

© 2026 IP Infusion Inc. Proprietary

	area).
(config-router)#exit	Exit from router ospf mode
(config)# router ldp	Configure Router LDP instance
(config-router)# transport-address ipv4 4.4.4.53 0	Configure Transport address for LDP with label space value 0
(config-router)#exit	Exit from router mode
(config)#commit	Commit the candidate configuration to the running configuration.

ABR2 Configuration

The following are the ABR2 configuration procedures.

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 5.5.5.56/32	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)#interface eth1	Enter interface mode
(config-if)#ip address 50.1.1.56/24	Configure the IP address of the interface eth1
(config-if)#label-switching	Enable label-switching on interface eth1
(config-if)# enable-ldp ipv4	Enable LDP process on eth1 interface
(config-if)#exit	Exit interface mode.
(config)#interface eth2	Enter interface mode
(config-if)#ip address 40.1.1.56/24	Configure the IP address of the interface eth2
(config-if)#label-switching	Enable label-switching on interface eth2
(config-if)# enable-ldp ipv4	Enable LDP process on eth2 interface
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router ospf 30	Configure the routing process OSPF with process id 20
(config-router)#network 5.5.5.56/32 area 0	Define the interface 5.5.5.56/32) 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)#exit`	Exit from router ospf mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router ospf 20`	Configure the routing process OSPF with process id 20
`(config-router)#network 5.5.5.56/32 area 0`	Define the interface 5.5.5.56/32) 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)#exit`	Exit from router ospf mode
`(config)# router ldp`	Configure Router LDP instance
`(config-router)# transport-address ipv4 3.3.3.52 0`	Configure Transport address for LDP with label space value 0
`(config-router)#exit`	Exit from router mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router bgp 100`	Enter Router BGP mode
`(config-router)#neighbor 3.3.3.52 remote-as 100`	Add loopback ip of PE2 as neighbor with neighbor AS
`(config-router)#neighbor 3.3.3.52 update-source lo`	Update the source for that particular neighbor as loopback interface
`(config-router)#neighbor 7.7.7.55 remote-as 100`	Add loopback ip of ABR1 as neighbor with neighbor AS
`(config-router)#neighbor 7.7.7.55 update-source lo`	Update the source for that particular neighbor as loopback interface
`(config-router)#address-family ipv4 labeled-unicast`	Enter into labeled-unicast address family
`(config-router-af)#neighbor 3.3.3.52 activate`	Activate the neighbor inside labeled-unicast address family
`(config-router-af)#neighbor 7.7.7.55 activate`	Activate the neighbor inside labeled-unicast address family
`(config-router-af)# neighbor 3.3.3.52 route-reflector-client`	Enable Route reflector client for the neighbor inside address family IPv4 labeled unicast
`(config-router-af)# neighbor 7.7.7.55 route-reflector-client`	Enable Route reflector client for the neighbor inside address family IPv4 labeled unicast
`(config-router-af)# neighbor 3.3.3.52 next-hop-self`	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast

© 2026 IP Infusion Inc. Proprietary

(config-router-af)# neighbor 7.7.7.55 next-hop-self	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast
(config-router-af)#exit-address-family	Exit from address family labeled-unicast
(config-router)#exit	Exit from router BGP mode
(config)#commit	Commit the candidate configuration to the running configuration.

P3 Configuration

The following are the P3 configuration procedures.

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 6.6.6.22/32	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)#interface eth1	Enter interface mode
(config-if)#ip address 50.1.1.22/24	Configure the IP address of the interface eth1
(config-if)#label-switching	Enable label-switching on interface eth1
(config-if)# enable-ldp ipv4	Enable LDP process on eth1 interface
(config-if)#exit	Exit interface mode.
(config)#interface eth2	Enter interface mode
(config-if)#ip address 60.1.1.22/24	Configure the IP address of the interface eth2
(config-if)#label-switching	Enable label-switching on interface eth2
(config-if)# enable-ldp ipv4	Enable LDP process on eth2 interface
(config-if)#exit	Exit interface mode
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router ospf 30	Configure the routing process OSPF with process id 20
(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 (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)#network 6.6.6.22/32 area 0	Define the interface (4.4.4.53/32) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone

© 2026 IP Infusion Inc. Proprietary

	area).
(config-router)#exit	Exit from router ospf mode
(config)#commit	Commit the candidate configuration to the running configuration.
(config)# router ldp	Configure Router LDP instance
(config-router)# transport-address ipv4 6.6.6.22 0	Configure Transport address for LDP with label space value 0
(config-router)#exit	Exit from router mode
(config)# ip route 1.1.1.54/32 50.1.1.56	Specify the destination prefix and mask for the network and a gateway.
(config)#exit	Exit from config mode
(config)#commit	Commit the candidate configuration to the running configuration.

PE2 Configuration

The following are the PE2 configuration procedures.

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 7.7.7.55/32 secondary	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)#interface eth2	Enter interface mode
(config-if)#ip address 60.1.1.55/24	Configure the IP address of the interface eth1
(config-if)#label-switching	Enable label-switching on interface eth1
(config-if)# enable-ldp ipv4	Enable LDP process on eth1 interface
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router ospf 30	Configure the routing process OSPF with process id 10
(config-router)#network 7.7.7.55/32 area 0	Define the interface (7.7.7.55/32) 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).

© 2026 IP Infusion Inc. Proprietary

`(config-router)#exit`	Exit from router ospf mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router bgp 100`	Enter Router BGP mode
`(config-router)#neighbor 1.1.1.54 remote-as 100`	Add loopback ip of PE2 as neighbor with neighbor AS
`(config-router)#neighbor 1.1.1.54 update-source lo`	Update the source for that particular neighbor as loopback interface
`(config-router)#neighbor 5.5.5.56 remote-as 100`	Add loopback ip of ABR1 as neighbor with neighbor AS
`(config-router)#neighbor 5.5.5.56 update-source lo`	Update the source for that particular neighbor as loopback interface
`(config-router)#allocate-label all`	Allocate labels
`(config-router)#address-family ipv4 labeled-unicast`	Enter into labeled-unicast address family
`(config-router-af)#neighbor 5.5.5.56 activate`	Activate the neighbor inside labeled-unicast address family
`(config-router-af)#exit-address-family`	Exit from address family IPv4 labeled unicast
`(config-router)#address-family vpnv4 unicast`	Enter into vpnv4 unicast address family
`(config-router-af)#neighbor 1.1.1.54 activate`	Activate the neighbor inside vpnv4 address family
`(config-router-af)#exit-address-family`	Exit from address family vpnv4
`(config-router)#address-family ipv4 unicast`	Enter into ipv4 unicast address family
`(config-router-af)#network 7.7.7.55/32`	Advertise the loopback of RTR1 in BGP
`(config-router-af)#exit-address-family`	Exit from address family
`(config-router)#exit`	Exit from router BGP mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)# router ldp`	Configure Router LDP instance
`(config-router)# transport-address ipv4 7.7.7.55 0`	Configure Transport address for LDP with label space value 0
`(config-router)#exit`	Exit from router mode
`(config)# ip vrf vrf1`	Specify the name of the VRF (vrf1) 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 100:300`	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 eth2	Enter interface mode
(config-if)# ip vrf forwarding vrf1	Bind the interface (eth2) to the VRF vrf1
(config-if)# ip address 62.1.1.55/24	Configure the IP address of the interface eth2
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router bgp 100	Enter Router BGP mode
(config-router)# address-family ipv4 vrf vrf1	Enter address family ipv4 vrf mode
(config-router-af)# redistribute connected	Redistribute connected routes
(config-router-af)# exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

BGP LU with Seamless MPLS Validation

PE1

```
#show ip bgp neighbors 3.3.3.52
BGP neighbor is 3.3.3.52, remote AS 100, local AS 100, internal link
 BGP version 4, local router ID 1.1.1.54, remote router ID 3.3.3.52
 BGP state = Established, up for 00:00:06
 Last read 00:00:07, hold time is 90, keepalive interval is 30 seconds Neighbor capabilities:
 Route refresh: advertised and received (old and new)
 Received 4 messages, 1 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 Update source is lo
 For address family: IPv4 Unicast
 BGP table version 1, neighbor version 1 Index 0, Offset 0, Mask 0x1
 Community attribute sent to this neighbor (both)
 0   accepted prefixes
 1   announced prefixes

For address family: IPv4 Labeled-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 2; dropped 1
Local host: 1.1.1.54, Local port: 179
Foreign host: 3.3.3.52, Foreign port: 46745
Nexthop: 1.1.1.54
Nexthop global: 54::54 Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:00:11, due to BGP Notification received Notification Error Message: (Cease/Other
Configuration Change.)

#show ip bgp neighbors 7.7.7.55
BGP neighbor is 7.7.7.55, remote AS 100, local AS 100, internal link
BGP version 4, local router ID 1.1.1.54, remote router ID 7.7.7.55
BGP state = Established, up for 00:01:10
Last read 00:00:11, hold time is 90, keepalive interval is 30 seconds Neighbor capabilities:
Route refresh: advertised and received (old and new)
Received 8 messages, 1 notifications, 0 in queue
```

```
Sent 9 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: VPNv4 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 3; dropped 2
Local host: 1.1.1.54, Local port: 179


Foreign host: 7.7.7.55, Foreign port: 58871
Nexthop: 1.1.1.54
Nexthop global: 54::54 Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:01:10, due to BGP Notification sent Notification Error Message:
(Cease/Administratively Reset.)


   #show mpls vrf-table
Output for IPv4 VRF table with id: 2
Primary FTN entry with FEC: 62.1.1.0/24, id: 1, row status: Active
Owner: BGP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none Tunnel id:
0,   Protected LSP id: 0, Description: N/A
Cross connect ix: 12, in intf: - in label: 0 out-segment ix: 11 Owner: BGP, Persistent: No, Admin
Status: Up, Oper Status: Up
Out-segment with ix: 11, owner: BGP, out intf: eth1, out label: 24961 Nexthop addr: 7.7.7.55   cross
connect ix: 12, op code: Push and
Lookup

#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
Network   Next Hop   In Label   Out Label
*>   1.1.1.54/32   0.0.0.0   24961   -
*>i 7.7.7.55/32   3.3.3.52   -   24961
#show mpls forwarding-table | include 7.7.7.55
B>   7.7.7.55/32   4   0   Yes   LSP_DEFAULT   24961
eth1   3.3.3.52

#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM K - CLI ILM, T - MPLS-TP, S - Stitched ILM

Code   FEC   ILM-ID   In-Label   Out-Label   In-Intf   Out- Intf      Nexthop         LSP-Type
>   1.1.1.54/32   11   24961   N/A   N/A   N/A
127.0.0.1   LSP_DEFAULT
>   61.1.1.0/24   13   24963   N/A   N/A   eth2
0.0.0.0   LSP_DEFAULT

#show mpls ftn-table
Primary FTN entry with FEC: 2.2.2.23/32, id: 1, row status: Active
Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none Tunnel id:
0,   Protected LSP id: 0, Description: N/A
Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1 Owner: LDP, Persistent: No, Admin
Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: LDP, out intf: eth1, out label: 3 Nexthop addr: 10.1.1.23   cross
connect ix: 1, op code: Push

Primary FTN entry with FEC: 3.3.3.52/32, id: 2, row status: Active
Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none Tunnel id:
0,   Protected LSP id: 0, Description: N/A
Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2 Owner: LDP, Persistent: No, Admin
Status: Up, Oper Status: Up
Out-segment with ix: 2, owner: LDP, out intf: eth1, out label: 24320


Nexthop addr: 10.1.1.23   cross connect ix: 2, op code: Push
```

```
Primary FTN entry with FEC: 7.7.7.55/32, id: 4, row status: Active
Owner: BGP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none Tunnel id:
0,   Protected LSP id: 0, Description: N/A
Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3 Owner: BGP, Persistent: No, Admin
Status: Up, Oper Status: Up
Out-segment with ix: 3, owner: BGP, out intf: eth1, out label: 24961 Nexthop addr: 3.3.3.52   cross
connect ix: 4, op code: Push and
Lookup


Primary FTN entry with FEC: 20.1.1.0/24, id: 3, row status: Active
Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none Tunnel id:
0,   Protected LSP id: 0, Description: N/A
Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1 Owner: LDP, Persistent: No, Admin
Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: LDP, out intf: eth1, out label: 3 Nexthop addr: 10.1.1.23   cross
connect ix: 1, op code: Push
```

P1

```
#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 10 VRF(default):
Neighbor ID     Pri   State            Dead Time   Address         Interface            Instance ID
1.1.1.54          1   Full/Backup      00:00:38    10.1.1.54       eth1                     0
3.3.3.52          1   Full/DR          00:00:39    20.1.1.52       eth2                     0
#show ldp session
Peer IP Address           IF Name    My Role    State      KeepAlive  UpTime
1.1.1.54                  eth1       Active    OPERATIONAL   30    00:00:21
3.3.3.52                  eth2       Passive   OPERATIONAL   30    00:00:15
```

ABR1

```
#show ip bgp neighbors 1.1.1.54
BGP neighbor is 1.1.1.54, remote AS 100, local AS 100, internal link BGP version 4, local router ID
3.3.3.52, remote router ID 1.1.1.54 BGP state = Established, up for 00:00:09
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 Address family IPv4 Labeled-Unicast: advertised and received
Received 5 messages, 0 notifications, 0 in queue
Sent 4 messages, 1 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 0, Offset 0, Mask 0x1
Community attribute sent to this neighbor (both)
1 accepted prefixes


0   announced prefixes

For address family: IPv4 Labeled-UnicastBGP table version 2, neighbor version 2 Index 0, Offset 0,
Mask 0x1
Route-Reflector Client NEXT_HOP is always this router
Community attribute sent to this neighbor (both)
1   accepted prefixes
0   announced prefixes

Connections established 2; dropped 1
Local host: 3.3.3.52, Local port: 46745
Foreign host: 1.1.1.54, Foreign port: 179
Nexthop: 3.3.3.52
Nexthop global: 52::52 Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:00:14, due to BGP Notification sent Notification Error Message: (Cease/Other
```

```
Configuration Change.)

#show ip bgp neighbors 5.5.5.56
BGP neighbor is 5.5.5.56, remote AS 100, local AS 100, internal link BGP version 4, local router ID
3.3.3.52, remote router ID 5.5.5.56 BGP state = Established, up for 00:00:11
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 IPv4 Labeled-Unicast: advertised and received
Received 4 messages, 0 notifications, 0 in queue
Sent 6 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)
1 accepted prefixes
1 announced prefixes

For address family: IPv4 Labeled-UnicastBGP table version 3, neighbor version 2 Index 1, Offset 0,
Mask 0x2
Route-Reflector Client NEXT_HOP is always this router
Community attribute sent to this neighbor (both)
1 accepted prefixes
1 announced prefixes

Connections established 2; dropped 1
Local host: 3.3.3.52, Local port: 179
Foreign host: 5.5.5.56, Foreign port: 40440
Nexthop: 3.3.3.52
Nexthop global: 52::52 Nexthop local: ::


BGP connection: non shared network
Last Reset: 00:00:11, due to BGP Notification sent Notification Error Message:
(Cease/Administratively Reset.)

#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.54/32       1.1.1.54        24960        24961
*>i 7.7.7.55/32       5.5.5.56        24961        24961

#show mpls forwarding-table | include 1.1.1.54
   L>    1.1.1.54/32         1        0           Yes   LSP_
DEFAULT     24321        eth2          20.1.1.23
   B     1.1.1.54/32         7        0           Yes   LSP_
DEFAULT     24961        eth2          1.1.1.54

#show mpls forwarding-table | include 7.7.7.55
   B>    7.7.7.55/32         8        0           Yes   LSP_
DEFAULT     24961        eth1          5.5.5.56

#show mpls ftn-table
 Primary FTN entry with FEC: 1.1.1.54/32, id: 1, row status: Active
  Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 2, owner: LDP, out intf: eth2, out label: 24321
    Nexthop addr: 20.1.1.23        cross connect ix: 2, op code: Push


 Primary FTN entry with FEC: 1.1.1.54/32, id: 7, row status: Active
  Owner: BGP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 10
      Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
       Out-segment with ix: 10, owner: BGP, out intf: eth2, out label: 24961
```

```
   Nexthop addr: 1.1.1.54        cross connect ix: 11, op code: Push and Lookup


Primary FTN entry with FEC: 2.2.2.23/32, id: 2, row status: Active
 Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 4
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 4, owner: LDP, out intf: eth2, out label: 3
   Nexthop addr: 20.1.1.23        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 4.4.4.53/32, id: 4, row status: Active
 Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, out intf: eth1, out label: 3
   Nexthop addr: 30.1.1.53        cross connect ix: 7, op code: Push


Primary FTN entry with FEC: 5.5.5.56/32, id: 5, row status: Active
 Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 9
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 9, owner: LDP, out intf: eth1, out label: 24320
   Nexthop addr: 30.1.1.53        cross connect ix: 10, op code: Push


Primary FTN entry with FEC: 7.7.7.55/32, id: 8, row status: Active
 Owner: BGP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Cross connect ix: 12, in intf: - in label: 0 out-segment ix: 11
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 11, owner: BGP, out intf: eth1, out label: 24961
   Nexthop addr: 5.5.5.56        cross connect ix: 12, op code: Push and Lookup


Primary FTN entry with FEC: 10.1.1.0/24, id: 3, row status: Active
 Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 4
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 4, owner: LDP, out intf: eth2, out label: 3
   Nexthop addr: 20.1.1.23        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 40.1.1.0/24, id: 6, row status: Active
 Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, out intf: eth1, out label: 3
   Nexthop addr: 30.1.1.53        cross connect ix: 7, op code: Push
```

P2

```
#show ip ospf neighbor

Total number of full neighbors: 2 OSPF process 20 VRF(default):
Neighbor ID   Pri   State   Dead Time   Address   Interface Instance ID
3.3.3.52   1   Full/Backup   00:00:39   30.1.1.52   eth1
0
5.5.5.56   1   Full/DR   00:00:40   40.1.1.56   eth2
0
```

```
#show ldp session
Peer IP Address   IF Name   My Role   State   KeepAlive   UpTime
3.3.3.52   eth1   Active   OPERATIONAL   30   00:00:21
5.5.5.56   eth2   Passive   OPERATIONAL   30   00:00:19
```

ABR2

```
#show ip bgp neighbors 3.3.3.52


BGP neighbor is 3.3.3.52, remote AS 100, local AS 100, internal link BGP version 4, local router ID
5.5.5.56, remote router ID 3.3.3.52 BGP state = Established, up for 00:00:12
Last read 00:00:07, 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 IPv4 Labeled-Unicast: advertised and received
Received 5 messages, 1 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 Update source is lo
For address family: IPv4 Unicast
BGP table version 1, neighbor version 1 Index 0, Offset 0, Mask 0x1
Community attribute sent to this neighbor (both)
1 accepted prefixes
1 announced prefixes

For address family: IPv4 Labeled-UnicastBGP table version 2, neighbor version 2 Index 0, Offset 0,
Mask 0x1
Route-Reflector Client NEXT_HOP is always this router
Community attribute sent to this neighbor (both)
1 accepted prefixes
1 announced prefixes

Connections established 2; dropped 1
Local host: 5.5.5.56, Local port: 40440
Foreign host: 3.3.3.52, Foreign port: 179
Nexthop: 5.5.5.56
Nexthop global: 56::56 Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:00:17, due to BGP Notification received Notification Error Message: (Cease/Other
Configuration Change.)

#show ip bgp neighbors 7.7.7.55
BGP neighbor is 7.7.7.55, remote AS 100, local AS 100, internal link BGP version 4, local router ID
5.5.5.56, remote router ID 7.7.7.55 BGP state = Established, up for 00:00:13
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 Address family IPv4 Labeled-Unicast: advertised and received
Received 3 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 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)


1 accepted prefixes
1 announced prefixes

For address family: IPv4 Labeled-UnicastBGP table version 2, neighbor version 2 Index 1, Offset 0,
Mask 0x2
Route-Reflector Client NEXT_HOP is always this router
Community attribute sent to this neighbor (both)
1 accepted prefixes
1 announced prefixes

Connections established 1; dropped 0
```

```
Local host: 5.5.5.56, Local port: 35004
Foreign host: 7.7.7.55, Foreign port: 179
Nexthop: 5.5.5.56
Nexthop global: 56::56 Nexthop local: ::
BGP connection: non shared network

#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.54/32       3.3.3.52        24960        24960
*>i 7.7.7.55/32       7.7.7.55        24961        24960

#show mpls forwarding-table | include 1.1.1.54
   B>    1.1.1.54/32         7        0           Yes   LSP_
DEFAULT     24960        eth2          3.3.3.52

#show mpls forwarding-table | include 7.7.7.55
   L>    7.7.7.55/32         5        0           Yes   LSP_
DEFAULT     24320        eth1           50.1.1.22
   B     7.7.7.55/32         8        0           Yes   LSP_
DEFAULT     24960        eth1           7.7.7.55

#show mpls ftn-table
 Primary FTN entry with FEC: 1.1.1.54/32, id: 7, row status: Active
  Owner: BGP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 10
      Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 10, owner: BGP, out intf: eth2, out label: 24960
    Nexthop addr: 3.3.3.52        cross connect ix: 11, op code: Push and Lookup


 Primary FTN entry with FEC: 3.3.3.52/32, id: 1, row status: Active
  Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 3, owner: LDP, out intf: eth2, out label: 24321
    Nexthop addr: 40.1.1.53        cross connect ix: 4, op code: Push


 Primary FTN entry with FEC: 4.4.4.53/32, id: 2, row status: Active
  Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 4
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 4, owner: LDP, out intf: eth2, out label: 3
    Nexthop addr: 40.1.1.53        cross connect ix: 5, op code: Push


 Primary FTN entry with FEC: 6.6.6.22/32, id: 4, row status: Active
  Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 6
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 6, owner: LDP, out intf: eth1, out label: 3
    Nexthop addr: 50.1.1.22        cross connect ix: 7, op code: Push


 Primary FTN entry with FEC: 7.7.7.55/32, id: 5, row status: Active
  Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 9
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 9, owner: LDP, out intf: eth1, out label: 24320
    Nexthop addr: 50.1.1.22        cross connect ix: 10, op code: Push
```

```
Primary FTN entry with FEC: 7.7.7.55/32, id: 8, row status: Active
 Owner: BGP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Cross connect ix: 12, in intf: - in label: 0 out-segment ix: 11
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 11, owner: BGP, out intf: eth1, out label: 24960
   Nexthop addr: 7.7.7.55        cross connect ix: 12, op code: Push and Lookup


Primary FTN entry with FEC: 30.1.1.0/24, id: 3, row status: Active
 Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 4
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 4, owner: LDP, out intf: eth2, out label: 3
   Nexthop addr: 40.1.1.53        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 60.1.1.0/24, id: 6, row status: Active
 Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, out intf: eth1, out label: 3
   Nexthop addr: 50.1.1.22        cross connect ix: 7, op code: Push
```

P3

```
#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 30 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface            Instance ID
5.5.5.56          1   Full/Backup     00:00:39    50.1.1.56       eth1                      0
7.7.7.55          1   Full/DR         00:00:39    60.1.1.55       eth2                     0

#show ldp session
Peer IP Address           IF Name    My Role    State      KeepAlive  UpTime
5.5.5.56                  eth1        Active    OPERATIONAL   30    00:00:30
7.7.7.55                  eth2       Passive   OPERATIONAL   30    00:00:28
```

PE2

```
#show ip bgp neighbors 5.5.5.56
BGP neighbor is 5.5.5.56, remote AS 100, local AS 100, internal link BGP version 4, local router ID
7.7.7.55, remote router ID 5.5.5.56 BGP state = Established, up for 00:00:16
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 Address family IPv4 Labeled-Unicast: advertised and received
Received 3 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 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)
1 accepted prefixes
1 announced prefixes

For address family: IPv4 Labeled-UnicastBGP table version 2, neighbor version 2 Index 1, Offset 0,
Mask 0x2
Community attribute sent to this neighbor (both)
1 accepted prefixes
1 announced prefixes
```

© 2026 IP Infusion Inc. Proprietary

```
Connections established 1; dropped 0
Local host: 7.7.7.55, Local port: 179
Foreign host: 5.5.5.56, Foreign port: 35004
Nexthop: 7.7.7.55
Nexthop global: 55::55 Nexthop local: ::
BGP connection: non shared network

#show mpls vrf-table
Output for IPv4 VRF table with id: 2
Primary FTN entry with FEC: 61.1.1.0/24, id: 1, row status: Active
Owner: BGP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none Tunnel id:
0,   Protected LSP id: 0, Description: N/A
Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 10


Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 10, owner: BGP, out intf: eth2, out label: 24963 Nexthop addr: 1.1.1.54   cross
connect ix: 11, op code: Push and
Lookup

#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
Network   Next Hop   In Label   Out Label
*>i 1.1.1.54/32   5.5.5.56   -   24960
*>   7.7.7.55/32   0.0.0.0   24960   -
#show mpls forwarding-table | include 1.1.1.54
B>   1.1.1.54/32   4   0   Yes   LSP_DEFAULT   24960
eth2   5.5.5.56

#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM K - CLI ILM, T - MPLS-TP, S - Stitched ILM

Code   FEC   ILM-ID   In-Label   Out-Label   In-Intf   Out- Intf      Nexthop         LSP-Type
>   7.7.7.55/32   9   24960   N/A   N/A   N/A
127.0.0.1   LSP_DEFAULT
>   62.1.1.0/24   10   24961   N/A   N/A   eth1
0.0.0.0   LSP_DEFAULT

#show mpls ftn-table
Primary FTN entry with FEC: 1.1.1.54/32, id: 4, row status: Active
Owner: BGP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none Tunnel id:
0,   Protected LSP id: 0, Description: N/A
Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3 Owner: BGP, Persistent: No, Admin
Status: Up, Oper Status: Up
Out-segment with ix: 3, owner: BGP, out intf: eth2, out label: 24960 Nexthop addr: 5.5.5.56   cross
connect ix: 4, op code: Push and
Lookup

Primary FTN entry with FEC: 5.5.5.56/32, id: 1, row status: Active
Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none Tunnel id:
0,   Protected LSP id: 0, Description: N/A
Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1 Owner: LDP, Persistent: No, Admin
Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: LDP, out intf: eth2, out label: 24321 Nexthop addr: 60.1.1.22   cross
connect ix: 1, op code: Push

Primary FTN entry with FEC: 6.6.6.22/32, id: 2, row status: Active
Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none Tunnel id:
0,   Protected LSP id: 0, Description: N/A
Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2 Owner: LDP, Persistent: No, Admin
Status: Up, Oper Status: Up
Out-segment with ix: 2, owner: LDP, out intf: eth2, out label: 3 Nexthop addr: 60.1.1.22   cross
connect ix: 2, op code: Push

Primary FTN entry with FEC: 50.1.1.0/24, id: 3, row status: Active


Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none Tunnel id:
0,   Protected LSP id: 0, Description: N/A
```

```
Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2 Owner: LDP, Persistent: No, Admin
Status: Up, Oper Status: Up
Out-segment with ix: 2, owner: LDP, out intf: eth2, out label: 3 Nexthop addr: 60.1.1.22   cross
connect ix: 2, op code: Push
```

BGP Labeled Unicast with Inter-AS Configuration

Topology

Figure 78. BGP Labeled unicast with Inter-AS

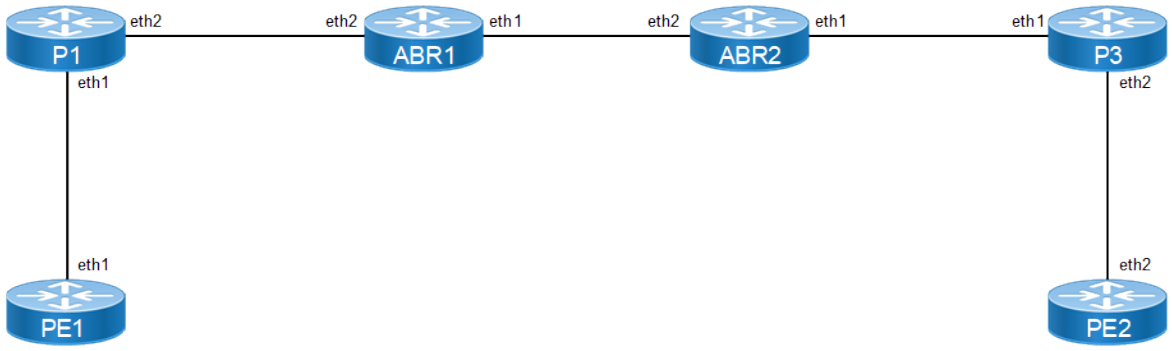

PE1 Configuration

The following are the PE1 configuration procedures.

`#configure terminal`	Enter the Configure mode.
`(config)#interface lo`	Enter interface mode
`(config-if)#ip address 1.1.1.54/32 secondary`	Configure the IP address of the interface loopback
`(config-if)#exit`	Exit interface mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip address 10.1.1.54/24`	Configure the IP address of the interface eth1
`(config-if)#label-switching`	Enable label-switching on interface eth1
`(config-if)# enable-ldp ipv4`	Enable LDP process on eth1 interface
`(config-if)#exit`	Exit interface mode.
`(config)#router ospf 10`	Configure the routing process OSPF with process id 10
`(config-router)#network 1.1.1.54/32 area 0`	Define the interface (1.1.1.54/32) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#network 10.1.1.0/24 area 0`	Define the interface (10.1.1.0/24) on which OSPF runs, and associate the area ID (0) with

© 2026 IP Infusion Inc. Proprietary

	the interface (area ID 0 specifies the backbone area).
`(config-router)#exit`	Exit from router ospf mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router bgp 100`	Enter Router BGP mode
`(config-router)#neighbor 7.7.7.55 remote-as 100`	Add loopback ip of PE2 as neighbor with neighbor AS
`(config-router)#neighbor 7.7.7.55 update-source lo`	Update the source for that particular neighbor as loopback interface
`(config-router)#neighbor 3.3.3.52 remote-as 100`	Add loopback ip of ABR1 as neighbor with neighbor AS
`(config-router)#neighbor 3.3.3.52 update-source lo`	Update the source for that particular neighbor as loopback interface
`(config-router)#allocate-label all`	Allocate labels
`(config-router)#address-family ipv4 labeled-unicast`	Enter into labeled-unicast address family
`(config-router-af)#neighbor 3.3.3.52 activate`	Activate the neighbor inside labeled-unicast address family
`(config-router-af)#exit-address-family`	Exit from address family IPv4 labeled unicast
`(config-router)#address-family vpnv4 unicast`	Enter into vpnv4 unicast address family
`(config-router-af)#neighbor 7.7.7.55 activate`	Activate the neighbor inside vpnv4 address family
`(config-router-af)#network 1.1.1.54/32`	Advertise the loopback of RTR1 in BGP
`(config-router-af)#exit-address-family`	Exit from address family vpnv4
`(config-router)#exit`	Exit from router BGP mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)# router ldp`	Configure Router LDP instance
`(config-router)# transport-address ipv4 1.1.1.54 0`	Configure Transport address for LDP with label space value 0
`(config-router)#exit`	Exit from router mode
`(config)# ip vrf vrf1`	Specify the name of the VRF (vrf1) 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 100:300`	Specify the 2-Octet AS specific or IPv4 specific Transitive Route-Target extended community attribute.

© 2026 IP Infusion Inc. Proprietary

(config-vrf)#exit	Exit VRF mode, and return to Configure mode.
(config)#interface eth2	Enter interface mode
(config-if)# ip vrf forwarding vrf1	Bind the interface (eth2) to the VRF vrf1
(config-if)# ip address 61.1.1.54/24	Configure the IP address of the interface eth2
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router bgp 100	Enter Router BGP mode
(config-router)# address-family ipv4 vrf vrf1	Enter address family ipv4 vrf mode
(config-router-af)# redistribute connected	Redistribute connected routes
(config-router-af)# exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

P1 Configuration

The following are the P1 configuration procedures.

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 2.2.2.23/32	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)#interface eth1	Enter interface mode
(config-if)#ip address 10.1.1.23/24	Configure the IP address of the interface eth1
(config-if)#label-switching	Enable label-switching on interface eth1
(config-if)# enable-ldp ipv4	Enable LDP process on eth1 interface
(config-if)#exit	Exit interface mode.
(config)#interface eth2	Enter interface mode
(config-if)#ip address 20.1.1.23/24	Configure the IP address of the interface eth2
(config-if)#label-switching	Enable label-switching on interface eth2
(config-if)# enable-ldp ipv4	Enable LDP process on eth2 interface
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router ospf 10	Configure the routing process OSPF with process id 10
(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

© 2026 IP Infusion Inc. Proprietary

	the interface (area ID 0 specifies the backbone area).
(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 2.2.2.23/32 area 0	Define the interface (2.2.2.23/32) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
(config-router)#exit	Exit from router ospf mode
(config)# router ldp	Configure Router LDP instance
(config-router)# transport-address ipv4 2.2.2.23 0	Configure Transport address for LDP with label space value 0
(config-router)#exit	Exit from router mode
(config)# ip route 7.7.7.55/32 20.1.1.52	Specify the destination prefix and mask for the network and a gateway.
(config)#commit	Commit the candidate configuration to the running configuration.

ABR1 Configuration

The following are the ABR1 configuration procedures.

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 3.3.3.52/32	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)#interface eth1	Enter interface mode
(config-if)#ip address 52.56.1.52/24	Configure the IP address of the interface eth1
(config-if)#label-switching	Enable label-switching on interface eth1
(config-if)# enable-ldp ipv4	Enable LDP process on eth1 interface
(config-if)#exit	Exit interface mode.
(config)#interface eth2	Enter interface mode
(config-if)#ip address 20.1.1.52/24	Configure the IP address of the interface eth2
(config-if)#label-switching	Enable label-switching on interface eth2
(config-if)# enable-ldp ipv4	Enable LDP process on eth2 interface
(config-if)#exit	Exit interface mode.
(config)#router ospf 10	Configure the routing process OSPF with process id 20

© 2026 IP Infusion Inc. Proprietary

`(config-router)#network 3.3.3.52/32 area 0`	Define the interface (3.3.3.52/32) 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).
`(config-router)#exit`	Exit from router ospf mode
`(config)# router ldp`	Configure Router LDP instance
`(config-router)# transport-address ipv4 3.3.3.52 0`	Configure Transport address for LDP with label space value 0
`(config-router)#exit`	Exit from router mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router bgp 100`	Enter Router BGP mode
`(config-router)#neighbor 1.1.1.54 remote-as 100`	Add loopback ip of PE2 as neighbor with neighbor AS
`(config-router)#neighbor 1.1.1.54 update-source lo`	Update the source for that particular neighbor as loopback interface
`(config-router)#neighbor 52.56.1.56 remote-as 100`	Add loopback ip of ABR1 as neighbor with neighbor AS
`(config-router)#allocate-label all`	Allocate labels
`(config-router)#address-family ipv4 labeled-unicast`	Enter into labeled-unicast address family
`(config-router-af)#neighbor 1.1.1.54 activate`	Activate the neighbor inside labeled-unicast address family
`(config-router-af)#neighbor 52.56.1.56 activate`	Activate the neighbor inside labeled-unicast address family
`(config-router-af)# neighbor 1.1.1.54 route-reflector-client`	Enable Route reflector client for the neighbor inside address family IPv4 labeled unicast
`(config-router-af)# neighbor 1.1.1.54 next-hop-self`	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast
`(config-router-af)# neighbor 52.56.1.56 next-hop-self`	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast
`(config-router-af)#exit-address-family`	Exit from address family labeled-unicast
`(config-router)#exit`	Exit from router BGP mode
`(config)#commit`	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

ABR2 Configuration

The following are the ABR2 configuration procedures.

`#configure terminal`	Enter the Configure mode.
`(config)#interface lo`	Enter interface mode
`(config-if)#ip address 5.5.5.56/32`	Configure the IP address of the interface loopback
`(config-if)#exit`	Exit interface mode
`(config)#interface eth1`	Enter interface mode
`(config-if)#ip address 50.1.1.56/24`	Configure the IP address of the interface eth1
`(config-if)#label-switching`	Enable label-switching on interface eth1
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode
`(config-if)#ip address 52.56.1.56/24`	Configure the IP address of the interface eth2
`(config-if)#label-switching`	Enable label-switching on interface eth2
`(config-if)#exit`	Exit interface mode.
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router ospf 30`	Configure the routing process OSPF with process Id 1
`(config-router)#network 5.5.5.56/32 area 0`	Define the interface (5.5.5.56/32) 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)#exit`	Exit from router ospf mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)# router ldp`	Configure Router LDP instance
`(config-router)# transport-address ipv4 5.5.5.56 0`	Configure Transport address for LDP with label space value 0
`(config-router)#exit`	Exit from router mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router bgp 200`	Enter Router BGP mode
`(config-router)#neighbor 52.56.1.52 remote-as 100`	Add neighbor peer ip and neighbor AS

© 2026 IP Infusion Inc. Proprietary

(config-router)#neighbor 7.7.7.55 remote-as 100	Add loopback ip of PE2 as neighbor with neighbor AS
(config-router)#neighbor 7.7.7.55 update-source lo	Update the source for that particular neighbor as loopback interface
(config-router)#allocate-label all	Allocate labels
(config-router)#address-family ipv4 labeled-unicast	Enter into labeled-unicast address family
(config-router-af)#neighbor 52.56.1.52 activate	Activate the neighbor inside labeled-unicast address family
(config-router-af)#neighbor 7.7.7.55 activate	Activate the neighbor inside labeled-unicast address family
(config-router-af)# neighbor 7.7.7.55 route-reflector-client	Enable Route reflector client for the neighbor inside address family IPv4 labeled unicast
(config-router-af)# neighbor 52.56.1.52 next-hop-self	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast
(config-router-af)# neighbor 7.7.7.55 next-hop-self	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast
(config-router-af)#exit-address-family	Exit from address family vpnv4
(config-router)#exit	Exit from router BGP mode
(config)#commit	Commit the candidate configuration to the running configuration.

P2 Configuration

The following are the P2 configuration procedures.

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 6.6.6.22/32	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)#interface eth1	Enter interface mode
(config-if)#ip address 50.1.1.22/24	Configure the IP address of the interface eth1
(config-if)#label-switching	Enable label-switching on interface eth1
(config-if)# enable-ldp ipv4	Enable LDP process on eth1 interface
(config-if)#exit	Exit interface mode.
(config)#interface eth2	Enter interface mode
(config-if)#ip address 60.1.1.22/24	Configure the IP address of the interface eth2
(config-if)#label-switching	Enable label-switching on interface eth2
(config-if)# enable-ldp ipv4	Enable LDP process on eth2 interface

© 2026 IP Infusion Inc. Proprietary

(config-if)#exit	Exit interface mode
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router ospf 30	Configure the routing process OSPF with process id 20
(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 (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)#network 6.6.6.22/32 area 0	Define the interface (4.4.4.53/32) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
(config-router)#exit	Exit from router ospf mode
(config)#commit	Commit the candidate configuration to the running configuration.
(config)# router ldp	Configure Router LDP instance
(config-router)# transport-address ipv4 6.6.6.22 0	Configure Transport address for LDP with label space value 0
(config-router)#exit	Exit from router mode
(config)# ip route 1.1.1.54/32 50.1.1.56	Specify the destination prefix and mask for the network and a gateway.
(config)#exit	Exit from Cofig mode
(config)#commit	Commit the candidate configuration to the running configuration.

PE2 Configuration

The following are the PE2 configuration procedures.

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 7.7.7.55/32 secondary	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)#interface eth2	Enter interface mode
(config-if)#ip address 60.1.1.55/24	Configure the IP address of the interface eth1

© 2026 IP Infusion Inc. Proprietary

(config-if)#label-switching	Enable label-switching on interface eth1
(config-if)# enable-ldp ipv4	Enable LDP process on eth1 interface
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router ospf 30	Configure the routing process OSPF with process id 10
(config-router)#network 7.7.7.55/32 area 0	Define the interface (7.7.7.55/32) 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)#exit	Exit from router ospf mode
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router bgp 200	Enter Router BGP mode
(config-router)#neighbor 1.1.1.54 remote-as 100	Add loopback ip of PE2 as neighbor with neighbor AS
(config-router)#neighbor 1.1.1.54 update-source lo	Update the source for that particular neighbor as loopback interface
(config-router)#neighbor 1.1.1.54 ebgp-multihop	Enable neighbor connection between two eBGP
(config-router)#neighbor 5.5.5.56 remote-as 200	Add loopback ip of ABR1 as neighbor with neighbor AS
(config-router)#neighbor 5.5.5.56 update-source lo	Update the source for that particular neighbor as loopback interface
(config-router)#allocate-label all	Allocate labels
(config-router)#address-family ipv4 labeled-unicast	Enter into labeled-unicast address family
(config-router-af)#neighbor 5.5.5.56 activate	Activate the neighbor inside vpnv4 address family
(config-router-af)#exit-address-family	Exit from address family IPv4 labeled unicast
(config-router)#address-family vpnv4 unicast	Enter into vpnv4 unicast address family
(config-router-af)#neighbor 1.1.1.54 allow-ebgp-vpn	Allow eBGP neighbor to be a vpn peer.
(config-router-af)#neighbor 1.1.1.54 activate	Activate the neighbor inside labeled-unicast address family
(config-router-af)#exit-address-family	Exit from address family vpnv4

© 2026 IP Infusion Inc. Proprietary

(config-router)#address-family ipv4 unicast	Enter into ipv4 unicast address family
(config-router-af)#network 7.7.7.55/32	Advertise the loopback of RTR1 in BGP
(config-router-af)#exit-address-family	Exit from address family vpnv4
(config-router)#exit	Exit from router BGP mode
(config)# router ldp	Configure Router LDP instance
(config-router)# transport-address ipv4 7.7.7.55 0	Configure Transport address for LDP with label space value 0
(config-router)#exit	Exit from router mode
(config)# ip vrf vrf1	Specify the name of the VRF (vrf1) 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 100:300	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 eth2	Enter interface mode
(config-if)# ip vrf forwarding vrf1	Bind the interface (eth2) to the VRF vrf1
(config-if)# ip address 62.1.1.55/24	Configure the IP address of the interface eth2
(config-if)#exit	Exit interface mode.
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router bgp 100	Enter Router BGP mode
(config-router)# address-family ipv4 vrf vrf1	Enter address family ipv4 vrf mode
(config-router-af)# redistribute connected	Redistribute connected routes
(config-router-af)# exit-address-family	Exit address-family mode.
(config-router)#commit	Commit the candidate configuration to the running configuration.

BGP LU with Inter-AS Validation

PE1

```
#sh ip bgp labeled-unicast
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
   Network    Next Hop    In Label   Out Label
*>   1.1.1.54/32    0.0.0.0     24969             -
*>i   7.7.7.55/32    3.3.3.52       -           24322


#sh mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, B - BGP FTN, K - CLI FTN, t - tunnel
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
```

```
10.1.1.23
10.1.1.23
3.3.3.52
10.1.1.23

#sh mpls ftn-table

Nexthop

Primary FTN entry with FEC: 2.2.2.23/32, id: 1, row status: Active
Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none Tunnel id:
0,   Protected LSP id: 0, Description: N/A
Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1 Owner: LDP, Persistent: No, Admin
Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: LDP, out intf: eth1, out label: 3 Nexthop addr: 10.1.1.23   cross
connect ix: 1, op code: Push

Primary FTN entry with FEC: 3.3.3.52/32, id: 3, row status: Active
Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none Tunnel id:
0,   Protected LSP id: 0, Description: N/A
Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2 Owner: LDP, Persistent: No, Admin
Status: Up, Oper Status: Up
Out-segment with ix: 2, owner: LDP, out intf: eth1, out label: 24321 Nexthop addr: 10.1.1.23   cross
connect ix: 2, op code: Push

Primary FTN entry with FEC: 7.7.7.55/32, id: 4, row status: Active
Owner: BGP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none Tunnel id:
0,   Protected LSP id: 0, Description: N/A
Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3 Owner: BGP, Persistent: No, Admin
Status: Up, Oper Status: Up
Out-segment with ix: 3, owner: BGP, out intf: eth1, out label: 24322 Nexthop addr: 3.3.3.52   cross
connect ix: 4, op code: Push and
Lookup

Primary FTN entry with FEC: 20.1.1.0/24, id: 2, row status: Active
Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none Tunnel id:
0,   Protected LSP id: 0, Description: N/A
Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1 Owner: LDP, Persistent: No, Admin
Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: LDP, out intf: eth1, out label: 3 Nexthop addr: 10.1.1.23   cross
connect ix: 1, op code: Push
```

ABR1

```
#sh ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.54/32       1.1.1.54        24323        24969
*>  7.7.7.55/32       52.56.1.56      24322        24325

#sh mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
       B - BGP FTN, K - CLI FTN, t - tunnel
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

Code    FEC                FTN-ID    Tunnel-id   Pri   LSP-Type        Out-Label    Out-
Intf       Nexthop
   L>    1.1.1.54/32         2        0           Yes   LSP_
DEFAULT     24322        eth2          20.1.1.23
   B     1.1.1.54/32         5        0           Yes   LSP_
DEFAULT     24969        eth2          1.1.1.54
   L>    2.2.2.23/32         3        0           Yes   LSP_
DEFAULT     3            eth2          20.1.1.23
   B>    7.7.7.55/32         1        0           Yes   LSP_
DEFAULT     24325        eth1          52.56.1.56
```

```
   L>    10.1.1.0/24         4        0           Yes   LSP_
DEFAULT     3            eth2          20.1.1.23

#sh mpls ftn-table
 Primary FTN entry with FEC: 1.1.1.54/32, id: 2, row status: Active
  Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 2, owner: LDP, out intf: eth2, out label: 24322
    Nexthop addr: 20.1.1.23        cross connect ix: 2, op code: Push


 Primary FTN entry with FEC: 1.1.1.54/32, id: 5, row status: Active
  Owner: BGP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 8
      Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
       Out-segment with ix: 8, owner: BGP, out intf: eth2, out label: 24969
    Nexthop addr: 1.1.1.54        cross connect ix: 8, op code: Push and Lookup


 Primary FTN entry with FEC: 2.2.2.23/32, id: 3, row status: Active
  Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 3, owner: LDP, out intf: eth2, out label: 3
    Nexthop addr: 20.1.1.23        cross connect ix: 3, op code: Push


 Primary FTN entry with FEC: 7.7.7.55/32, id: 1, row status: Active
  Owner: BGP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 7
      Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 7, owner: BGP, out intf: eth1, out label: 24325
    Nexthop addr: 52.56.1.56        cross connect ix: 7, op code: Push


 Primary FTN entry with FEC: 10.1.1.0/24, id: 4, row status: Active
  Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 3, owner: LDP, out intf: eth2, out label: 3
    Nexthop addr: 20.1.1.23        cross connect ix: 3, op code: Push
```

ABR2

```
#sh ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>  1.1.1.54/32       52.56.1.52      24322        24323
*>i 7.7.7.55/32       7.7.7.55        24325        24967

#sh mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
       B - BGP FTN, K - CLI FTN, t - tunnel
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

Code    FEC                FTN-ID    Tunnel-id   Pri   LSP-Type        Out-Label    Out-
Intf       Nexthop
   B>    1.1.1.54/32         1        0           Yes   LSP_
DEFAULT     24323        eth2          52.56.1.52
```

```
   L>    6.6.6.22/32         2        0           Yes   LSP_
DEFAULT     3            eth1           50.1.1.22
   L>    7.7.7.55/32         3        0           Yes   LSP_
DEFAULT     24322        eth1           50.1.1.22
   B     7.7.7.55/32         5        0           Yes   LSP_
DEFAULT     24967        eth1           7.7.7.55
   L>    60.1.1.0/24         4        0           Yes   LSP_
DEFAULT     3            eth1           50.1.1.22

#sh mpls ftn-table
 Primary FTN entry with FEC: 1.1.1.54/32, id: 1, row status: Active
  Owner: BGP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 8
      Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 8, owner: BGP, out intf: eth2, out label: 24323
    Nexthop addr: 52.56.1.52        cross connect ix: 8, op code: Push


 Primary FTN entry with FEC: 6.6.6.22/32, id: 2, row status: Active
  Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 2, owner: LDP, out intf: eth1, out label: 3
    Nexthop addr: 50.1.1.22        cross connect ix: 2, op code: Push


 Primary FTN entry with FEC: 7.7.7.55/32, id: 3, row status: Active
  Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 3, owner: LDP, out intf: eth1, out label: 24322
    Nexthop addr: 50.1.1.22        cross connect ix: 3, op code: Push


 Primary FTN entry with FEC: 7.7.7.55/32, id: 5, row status: Active
  Owner: BGP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 7
      Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
       Out-segment with ix: 7, owner: BGP, out intf: eth1, out label: 24967
    Nexthop addr: 7.7.7.55        cross connect ix: 7, op code: Push and Lookup


 Primary FTN entry with FEC: 60.1.1.0/24, id: 4, row status: Active
  Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 2, owner: LDP, out intf: eth1, out label: 3
    Nexthop addr: 50.1.1.22        cross connect ix: 2, op code: Push
```

PE2

```
#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
    Network          Next Hop        In Label     Out Label
*>i 1.1.1.54/32       5.5.5.56           -            24322
*>  7.7.7.55/32      0.0.0.0         24967          -

#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
      B - BGP FTN, K - CLI FTN, t - tunnel
```

```
      L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
      U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

Code    FEC                FTN-ID    Tunnel-id   Pri   LSP-Type        Out-Label    Out-
Intf       Nexthop
  B>    1.1.1.54/32          4        0           Yes   LSP_
DEFAULT     24322        eth2          5.5.5.56
  L>    5.5.5.56/32          3        0           Yes   LSP_
DEFAULT     24321        eth2          60.1.1.22
  L>    6.6.6.22/32          2        0           Yes   LSP_
DEFAULT     3            eth2          60.1.1.22
  L>    50.1.1.0/24          1        0           Yes   LSP_
DEFAULT     3            eth2          60.1.1.22

#sh mpls ftn-table
 Primary FTN entry with FEC: 1.1.1.54/32, id: 4, row status: Active
  Owner: BGP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 36, in intf: - in label: 0 out-segment ix: 37
      Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 37, owner: BGP, out intf: eth2, out label: 24322
    Nexthop addr: 5.5.5.56        cross connect ix: 36, op code: Push and Lookup


 Primary FTN entry with FEC: 5.5.5.56/32, id: 3, row status: Active
  Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 6
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 6, owner: LDP, out intf: eth2, out label: 24321
    Nexthop addr: 60.1.1.22        cross connect ix: 7, op code: Push


 Primary FTN entry with FEC: 6.6.6.22/32, id: 2, row status: Active
  Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 8
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 8, owner: LDP, out intf: eth2, out label: 3
    Nexthop addr: 60.1.1.22        cross connect ix: 9, op code: Push


 Primary FTN entry with FEC: 50.1.1.0/24, id: 1, row status: Active
  Owner: LDP, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
     Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 8
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 8, owner: LDP, out intf: eth2, out label: 3
    Nexthop addr: 60.1.1.22        cross connect ix: 9, op code: Push
```

© 2026 IP Infusion Inc. Proprietary

BGP Peer Groups for Address-Family IPv4 Labeled Unicast

Topology

Figure 79. BGP Peer-Groups with IPv4 Labeled-Unicast Members

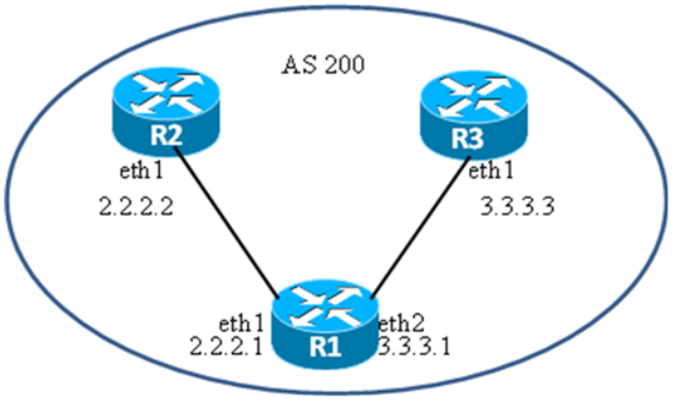

R1 Configuration

The following are the R1 configuration procedures.

`(config)# interface lo`	Enter interface mode for Loopback
`(config-if)#ip address 1.1.1.1/32 secondary`	Configure ip address for Loopback interface
`(config-if)#ip address 11.11.11.11/32 secondary`	Configure ip address for Loopback interface
`(config-if)#exit`	Exit interface mode
`(config)#router bgp 200`	Define the routing process. The number 200 specifies the AS number of R1.
`(config-router)#neighbor ABC peer-group`	Create a peer group named ABC.
`(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)#address-family ipv4 labeled-unicast`	Enter address-family ipv4 labeled-unicast mode
`(config-router-af)#neighbor ABC activate`	Activate the peer-group ABC for address-family ipv4 labeled-unicast
`(config-router-af)#neighbor ABC route-reflector-client`	Configure the peer-group ABC to be route-reflector-client
`(config-router-af)# exit-address-family`	Exit address-family ipv4 labeled-unicast mode
`(config-router)#address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)#network 1.1.1.1/32`	Advertise the network 1.1.1.1/32

© 2026 IP Infusion Inc. Proprietary

(config-router-af)#network 11.11.11.11/32	Advertise the network 11.11.11.11/32
(config-router-af)# exit-address-family	Exit address-family ipv4 unicast mode
(config-router)#allocate-label all	Allocate labels for all IPv4 prefixes advertised
(config-router)#exit	Exit router bgp mode
(config)#commit	Commit the candidate configuration to the running configuration.

R2 Configuration

The following are the R2 configuration procedures.

#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)#allocate-label all	Allocate labels for all IPv4 prefixes advertised
(config-router)#address-family ipv4 labeled-unicast	Enter address-family ipv4 labeled-unicast mode
(config-router-af)#neighbor 2.2.2.1 activate	Activate the neighbor R1 for address-family ipv4 labeled-unicast
(config-router-af)# exit-address-family	Exit address-family ipv4 labeled-unicast mode
(config-router)# exit	Exit router bgp mode
(config)#commit	Commit the candidate configuration to the running configuration.

R3 Configuration

The following are the R3 configuration procedures.

#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)#allocate-label all	Allocate labels for all IPv4 prefixes advertised
(config-router)#address-family ipv4 labeled-unicast	Enter address-family ipv4 labeled-unicast mode
(config-router-af)#neighbor 3.3.3.1 activate	Activate the neighbor R1 for address-family ipv4 labeled-unicast
(config-router-af)# exit-address-family	Exit address-family ipv4 labeled-unicast mode

© 2026 IP Infusion Inc. Proprietary

`(config-router)# exit`	Exit router bgp mode
`(config)#commit`	Commit the candidate configuration to the running configuration.

BGP Peer Groups for Address-Family IPv4 LU Validation

R1

```
R1#show ip bgp labeled-unicast summary
BGP router identifier 10.12.65.126, 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
2.2.2.2                  4   200   18         22       1      0      0  00:00:57               0
3.3.3.3                  4   200   18         20       1      0      0  00:00:01               0

Total number of neighbors 2

Total number of Established sessions 2
R1#


R1#show 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 10.12.65.126, remote router ID 10.12.65.123
  BGP state = Established, up for 00:01:05
  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
    Address family IPv4 Labeled-Unicast: advertised and received
  Received 16 messages, 2 notifications, 0 in queue
  Sent 20 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 11, neighbor version 11
  Index 0, Offset 0, Mask 0x1
  ABC peer-group member
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  2 announced prefixes

 For address family: IPv4 Labeled-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 5; dropped 4
Local host: 2.2.2.1, Local port: 51667
Foreign host: 2.2.2.2, Foreign port: 179
Nexthop: 2.2.2.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:01:10, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)
```

```
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 00:00:09
  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
    Address family IPv4 Labeled-Unicast: advertised and received
  Received 16 messages, 2 notifications, 0 in queue
  Sent 20 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 11, neighbor version 11
  Index 1, Offset 0, Mask 0x2
  ABC peer-group member
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  2 announced prefixes

 For address family: IPv4 Labeled-Unicast
  BGP table version 1, neighbor version 1
  Index 3, Offset 0, Mask 0x8
  ABC peer-group member
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  2 announced prefixes

 Connections established 5; dropped 4
Local host: 3.3.3.1, Local port: 41732
Foreign host: 3.3.3.3, Foreign port: 179
Nexthop: 3.3.3.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:00:19, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)
```

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
*>il 1.1.1.1/32       2.2.2.1               0         100           0        i
* i                   2.2.2.1               0         100           0        i
*>il 11.11.11.11/32   2.2.2.1               0         100           0        i
* i                   2.2.2.1               0         100           0        i

Total number of prefixes 2
R2#



R2#show ip bgp 1.1.1.1/32
BGP routing table entry for 1.1.1.1/32
Paths: (2 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
```

```
  Local
    2.2.2.1 from 2.2.2.1 (10.12.65.126)
      Origin IGP, metric 0, localpref 100      valid, internal, best, source safi: 4
      Last update: Tue Aug 25 10:01:05 2020

  Local
    2.2.2.1 from 2.2.2.1 (10.12.65.126)
      Origin IGP, metric 0, localpref 100      valid, internal, source safi: 1
      Last update: Tue Aug 25 10:01:05 2020

R2#
R2#show ip bgp 11.11.11.11/32
BGP routing table entry for 11.11.11.11/32
Paths: (2 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  Local
    2.2.2.1 from 2.2.2.1 (10.12.65.126)
      Origin IGP, metric 0, localpref 100      valid, internal, best, source safi: 4
      Last update: Tue Aug 25 10:01:05 2020

  Local
    2.2.2.1 from 2.2.2.1 (10.12.65.126)
      Origin IGP, metric 0, localpref 100      valid, internal, source safi: 1
      Last update: Tue Aug 25 10:01:05 2020

R2#


R2#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.1/32       2.2.2.1           -            24320
*>i 11.11.11.11/32   2.2.2.1           -            24321
R2#
```

R3

```
R3#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.1/32       3.3.3.1           -            24320
*>i 11.11.11.11/32   3.3.3.1           -            24321
R3#
```

Peer Group Outbound Policies Inheritance

Peer group members inherit the properties of Outbound Policies configured for Peer-group under Address-family IPv4 Labeled-Unicast.

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)#exit	Exit route-map mode
(config)#commit	Commit the candidate configuration to the running configuration.
(config)#router bgp 200	Define the routing process. The number 200 specifies the AS number of R1.
(config-router)#no allocate-label all	Unconfiguring allocate label all
config-router)# allocate-label route-map permit-only-1	Allocate labels to the route map
(config-router)#address-family ipv4 labeled-unicast	Enter address-family ipv4 labeled-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 address-family mode
(config-router)#exit	Exit router bgp mode
(config# exit	Exit configure terminal mode
(config)#commit	Commit the candidate configuration to the running configuration.
#clear ip bgp peer-group ABC ipv4 labeled- unicast soft out	Do outbound soft reset for the peer-group ABC for the policy to take affect for the labelled-unicast peer-group members

Validation

R1

```
R1#show 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 10.12.65.126, remote router ID 10.12.65.123
  BGP state = Established, up for 00:14:24
  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
    Address family IPv4 Labeled-Unicast: advertised and received
  Received 46 messages, 2 notifications, 0 in queue
  Sent 53 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 11, neighbor version 11
  Index 0, Offset 0, Mask 0x1
  ABC peer-group member
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  2 announced prefixes
```

```
 For address family: IPv4 Labeled-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 5; dropped 4
Local host: 2.2.2.1, Local port: 51667
Foreign host: 2.2.2.2, Foreign port: 179
Nexthop: 2.2.2.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:14:29, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)

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 00:13:28
  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
    Address family IPv4 Labeled-Unicast: advertised and received
  Received 47 messages, 2 notifications, 0 in queue
  Sent 53 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 11, neighbor version 11
  Index 1, Offset 0, Mask 0x2
  ABC peer-group member
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  2 announced prefixes

 For address family: IPv4 Labeled-Unicast
  BGP table version 1, neighbor version 1
  Index 3, Offset 0, Mask 0x8
  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 5; dropped 4
Local host: 3.3.3.1, Local port: 41732
Foreign host: 3.3.3.3, Foreign port: 179
Nexthop: 3.3.3.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:13:38, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)

R1#
```

R2

```
R2#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.1/32       2.2.2.1           -            24320
R2#
```

R3

```
R3#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.1/32       3.3.3.1           -            24320
R3#
```

Peer Group Inbound Policies Inheritance

Peer-group-members inherit the properties of Inbound Policies configured for Peer-group under Address-family ipv4 Labeled-Unicast.

Follow the configuration of R1, R2, R3 in the previous section with the following configuration

R2

The following are the R2 configuration procedures.

`#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 address-family mode
`(config-router)#allocate-label all`	Allocate labels for all IPv4 prefixes advertised
`(config-router)#exit`	Exit router bgp mode
`(config)#commit`	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

R3

The following are the R3 configuration procedures.

`#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-address-family`	Exit address-family mode
`(config-router)#allocate-label all`	Allocate labels for all IPv4 prefixes advertised
`(config-router)#exit`	Exit router bgp mode
`(config)#commit`	Commit the candidate configuration to the running configuration.

R1

The following are the R1 configuration procedures.

`#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)#commit`	Commit the candidate configuration to the running configuration.
`(config)#router bgp 200`	Enter BGP router mode
`(config-router)#address-family ipv4 labeled-unicast`	Enter the address-family ipv4 labeled-unicast mode
`(config-router-af)#neighbor ABC route-map permit-only-22 in`	Configure the peer-group ABC with route-map in the inbound direction under address-family ipv4 labeled-unicast

© 2026 IP Infusion Inc. Proprietary

`(config-router-af)# neighbor ABC activate`	Activate PEER-GROUP in the address family
`(config-router-af)#exit-address-family`	Exit address-family ipv4 labeled-unicast mode
`(config-router)#exit`	Exit router bgp mode
`(config)# exit`	Exit configure terminal mode
`(config)#commit`	Commit the candidate configuration to the running configuration.
`#clear ip bgp peer-group ABC ipv4 labeled-unicast soft in`	Do inbound soft reset for the peer-group ABC for the policy to take affect for the labelled-unicast peer-group members

Validation

R1

```
R1#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>  1.1.1.1/32       0.0.0.0          24320          -
*>  11.11.11.11/32   0.0.0.0          24321          -
*>i 22.1.1.0/24      3.3.3.3          24322        24320
* i                  2.2.2.2           -            24320
*>i 100.1.1.0/24     3.3.3.3          24323          -
R1#


R1#show 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 10.12.65.126, remote router ID 10.12.65.123
  BGP state = Established, up for 00:45:38
  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
    Address family IPv4 Labeled-Unicast: advertised and received
  Received 123 messages, 2 notifications, 0 in queue
  Sent 126 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 14, neighbor version 14
  Index 0, Offset 0, Mask 0x1
  ABC peer-group member
  Community attribute sent to this neighbor (both)
  2 accepted prefixes
  2 announced prefixes

 For address family: IPv4 Labeled-Unicast
  BGP table version 3, neighbor version 3
  Index 2, Offset 0, Mask 0x4
  ABC peer-group member
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Inbound path policy configured
  Outbound path policy configured
  Route map for incoming advertisements is *permit-only-22
  Route map for outgoing advertisements is *permit-only-1
  1 accepted prefixes
```

```
  1 announced prefixes

 Connections established 5; dropped 4
Local host: 2.2.2.1, Local port: 51667
Foreign host: 2.2.2.2, Foreign port: 179
Nexthop: 2.2.2.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:45:43, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)

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 00:44:42
  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
    Address family IPv4 Labeled-Unicast: advertised and received
  Received 124 messages, 2 notifications, 0 in queue
  Sent 127 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 14, neighbor version 14
  Index 1, Offset 0, Mask 0x2
  ABC peer-group member
  Community attribute sent to this neighbor (both)
  2 accepted prefixes
  2 announced prefixes

 For address family: IPv4 Labeled-Unicast
  BGP table version 3, neighbor version 3
  Index 3, Offset 0, Mask 0x8
  ABC peer-group member
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Inbound path policy configured
  Outbound path policy configured
  Route map for incoming advertisements is *permit-only-22
  Route map for outgoing advertisements is *permit-only-1
  1 accepted prefixes
  1 announced prefixes

 Connections established 5; dropped 4
Local host: 3.3.3.1, Local port: 41732
Foreign host: 3.3.3.3, Foreign port: 179
Nexthop: 3.3.3.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:44:52, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)
```

BGP Labeled Unicast with Implicit Null Label for Local Routes

Overview

In OcNOS, by default, the BGP router allocates a unique, valid MPLS label (>16) for all the IPv4 routes when `allocate-label (all | route-map NAME)` command is issued. The BGP Labeled Unicast (BGP LU) attaches the MPLS label to the routes and advertises them to its peers. This existing functionality is enhanced for Penultimate hop popping (PHP) requirement with an implicit null label allocation only for locally originated IPv4 routes, whether connected or redistributed from other protocols. Meanwhile, the non local routes are assigned with a valid MPLS label. Using PHP with implicit null eliminates the unnecessary label allocation, providing a more efficient and effective solution for handling local routes in the BGP-LU. If BGP LU is deployed in multiple segments, the implicit null label plays a significant role in the last segment.

A new CLI is introduced under the BGP router configuration mode. When the CLI is configured, the BGP implicit null label allocation is applied only after the expiry of the currently active BGP RIB scan timer that validates all the routes at the configured interval time.

This feature is supported on Qumran1 (Q1) series platforms, Qumran2 (Q2) series platforms, and Qumran (Q1 and Q2) series platforms (Q1, Q2 Qax Qux Q2c, Qmx).

Benefits

The regular MPLS label allocation for connected routes necessitates additional label processing on the egress Provider Edge (PE) and, sometimes, fails to forward traffic. The BGP LU with IMPLICIT NULL support reduces the label processing overheads at the egress node and efficiently uses the hardware FEC resources during route allocation.

Prerequisites

The following is mandatory before issuing the `bgp implicit-null` CLI:

- Set up the global BGP configuration using the command `allocate-label (all | route-map NAME)` to assign MPLS labels for IPv4 unicast routes. This configuration ensures that IPv4 unicast routes are qualified for MPLS label assignment. The following show running sample output illustrates that the MPLS labels are assigned to the routes.

```
!
router bgp 100
 allocate-label all
 neighbor 11.11.11.55 remote-as 100
 neighbor 172.3.4.53 remote-as 100
 neighbor 11.11.11.55 update-source lo
 !
 address-family ipv4 unicast
 network 21.21.21.56/32
 exit-address-family
 !
```

- Configuring BGP implicit null among network peers running different OcNOS versions earlier than 6.5.x causes traffic drops. To restore traffic flow, remove the implicit null configuration. Additionally, to implement this feature, ensure all peers are using version 6.5.x or higher.

© 2026 IP Infusion Inc. Proprietary

Topology

In this topology, PE1 node advertise all connected or local IPv4 unicast routes advertised in labeled-unicast address-family mode with implicit null label to remote PE2 via ABR node.

Figure 80. BGP labeled unicast with Implicit Null Label

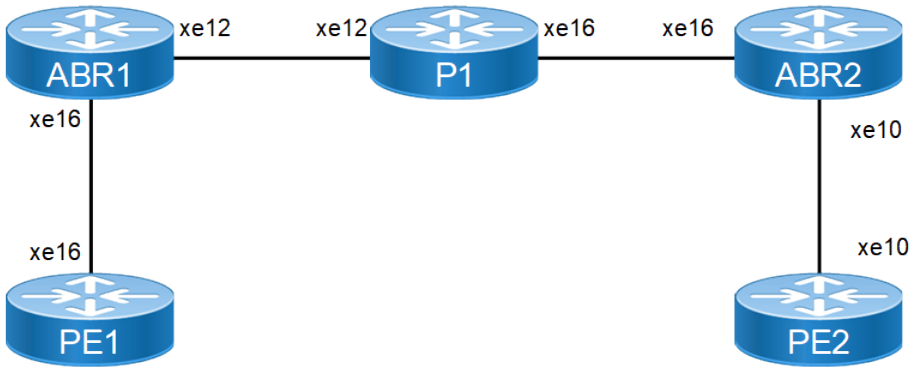

Configuration

The following configuration adds the implicit null label (3) to the locally connected IPv4 routes.

The BGP-LU with implicit null label introduces bgp implicit-null (page 1276) the configuration command.

BGP Labeled Unicast as Transport with Implicit Null Label

Perform the following configurations on OcNOS router to apply implicit null label.

1. Configure BGP-LU using the configuration provided in BGP Labeled Unicast (page 662) section.
2. In the BGP Labeled Unicast (page 662)onfiguration after the alocate-label all CLI issued in PE1 node, enable the `address-family IPv4 labeled-unicast` mode. Enabling the `address-family ipv4 labeled-unicast` mode, automatically adds bgp implicit-null CLI to the existing configuration.

```
(config-router)#allocate-label all
(config-router)#address-family ipv4 labeled- unicast
(config-router-af)#bgp implicit-null
```

> **Note:** Use of the `bgp implicit-null` command allows the BGP algorithm to automatically assign an implicit null label to local routes and a valid label (value > 16) to non-local routes.

Sample show running-config Output

The following is an example of the running configuration of PE1 device.

```
PE1#show running-config
feature netconf-ssh vrf management
feature netconf-tls vrf management
no feature netconf-ssh
no feature netconf-tls
no service password-encryption
!
```

```
logging console 7
logging monitor 7
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
!
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
hardware-profile port-config mode3
!
hostname PE1-7017
port ce46 breakout 4X10g
port cd50 breakout 4X100g
clock timezone kolkata
no ip domain-lookup
ip domain-lookup vrf management
ip name-server vrf management 10.12.3.23
ip name-server vrf management 10.12.3.25
ip name-server vrf management 10.12.3.35
bridge 1 protocol rstp vlan-bridge
tfo Disable
errdisable cause stp-bpdu-guard
no 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 community test group network-operator vrf management
snmp-server host 10.12.47.85 traps version 2c test udp-port 162 vrf management
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
ntp enable vrf management
ntp master vrf management
.
.
feature rsyslog vrf management
.
.
router bgp 100
 bgp router-id 192.168.1.1
 allocate-label all
 neighbor 192.168.1.4 remote-as 100
 neighbor 192.168.1.5 remote-as 100
 neighbor 192.168.1.6 remote-as 100
 neighbor 192.168.1.7 remote-as 100
 neighbor 192.168.1.4 update-source lo
 neighbor 192.168.1.4 advertisement-interval 0
 neighbor 192.168.1.5 update-source lo
 neighbor 192.168.1.5 advertisement-interval 0
 neighbor 192.168.1.6 update-source lo
 neighbor 192.168.1.6 advertisement-interval 0
 neighbor 192.168.1.7 update-source lo
 neighbor 192.168.1.7 advertisement-interval 0
 !
 address-family ipv4 unicast
 network 192.168.1.1/32
 exit-address-family
 !
 address-family ipv4 labeled-unicast
 bgp implicit-null
 neighbor 192.168.1.4 activate
 neighbor 192.168.1.5 activate
 exit-address-family
 !
 address-family vpnv4 unicast
```

© 2026 IP Infusion Inc. Proprietary

```
 neighbor 192.168.1.6 activate
 neighbor 192.168.1.7 activate
 exit-address-family
 !
 address-family ipv4 vrf vrf103
 redistribute connected
 exit-address-family
 !
 address-family ipv4 vrf vrf100
 redistribute connected
 exit-address-family
 !
 address-family ipv4 vrf vrf104
 redistribute connected
 exit-address-family
 !
 address-family ipv4 vrf vrf101
 redistribute connected
 exit-address-family
 !
 address-family ipv4 vrf vrf102
 redistribute connected
 exit-address-family
 !
 exit
!
.
.
end
```

Validation

Following output displays the BGP routes of neighbor configured on PE1.

PE1

```
PE1#show ip bgp neighbors 21.21.21.56
BGP neighbor is 21.21.21.56, remote AS 100, local AS 100, internal link, peer index: 3
  BGP version 4, local router ID 11.11.11.55, remote router ID 21.21.21.56
  BGP state = Established, up for 01:12:26
  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 VPNv4 Unicast: advertised and received
  Received 173 messages, 0 notifications, 0 in queue
  Sent 172 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: VPNv4 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
  1 announced prefixes

Connections established 1; dropped 0
Local host: 11.11.11.55, Local port: 35217
Foreign host: 21.21.21.56, Foreign port: 179
TCP MSS: (0), Advertise TCP MSS: (1460), Send TCP MSS: (1460),  Receive TCP MSS: (536)
Sock FD : (28)
Nexthop: 11.11.11.55
```

```
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network


PE1#show ip bgp neighbors 172.4.5.52
BGP neighbor is 172.4.5.52, remote AS 100, local AS 100, internal link, peer index: 4
  BGP version 4, local router ID 11.11.11.55, remote router ID 22.22.22.52
  BGP state = Established, up for 01:17:24
  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 Labeled-Unicast: advertised and received
  Received 273 messages, 1 notifications, 0 in queue
  Sent 272 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 Labeled-Unicast  BGP table version 4, neighbor version 4
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  1 accepted prefixes
  1 announced prefixes

Connections established 2; dropped 1
Local host: 172.4.5.55, Local port: 33740
Foreign host: 172.4.5.52, Foreign port: 179
TCP MSS: (0), Advertise TCP MSS: (1460), Send TCP MSS: (1460),  Receive TCP MSS: (536)
Sock FD : (27)
Nexthop: 172.4.5.55
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 01:17:29, due to BGP Notification received
Notification Error Message: (Cease/Administratively Reset.)
```

Following output displays the Labels attached to the BGP routes on PE1. The local route under “In Label” is replaced with 3.

```
PE1#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal, S - stale
     Network          Next Hop        In Label     Out Label
*>  11.11.11.55/32   0.0.0.0          3          -
*>i 21.21.21.56/32   172.4.5.52       25600        25601


PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled

Code    FEC                FTN-ID    Nhlfe-ID  Tunnel-id   Pri   LSP-Type     Out-Label    Out-
Intf    ELC   Nexthop
   B>   21.21.21.56/32      1         3         -           Yes   LSP_
DEFAULT  25601        ce46/1      No    172.4.5.52
PE1#

PE1#
PE1#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
```

```
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP , K - CLI , V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

ILM-ECMP LDP: Disabled

Code    FEC/VRF/L2CKT   ILM-ID      In-Label    Out-Label    In-Intf   Out-
Intf/VRF    Nexthop          pri    LSP-Type
   B>   vrf1              2          25601      Nolabel      N/A       N/A             N/A
Yes    LSP_DEFAULT
   B>   21.21.21.56/32    1          25600      25601        N/A       N/A             172.4.5.52
Yes    LSP_DEFAULT

PE1#
```

ABR1

Following output displays the BGP routes of neighbor configured on ABR1.

```
ABR1#show ip ospf neighbor

Total number of full neighbors: 2
      OSPF process 1 VRF(default):
      Neighbor ID     Pri   State            Dead Time   Address         Interface            Instance
ID
      11.11.11.55       1   Full/DR          00:00:34    172.4.5.55      xe3                     0
      200.200.200.54    1   Full/Backup      00:00:33    172.6.7.54      xe21                    0
```

Following output displays the Labels attached to the BGP routes on ABR1. The local route under “Out Label” is replaced with 3.

```
ABR1#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 11.11.11.55/32   172.4.5.55        25600        3
*>i 21.21.21.56/32   172.6.7.54        25601        25601

ABR1#show mpls forwarding-table | include 11.11.11.55

   B>   11.11.11.55/32      1         4          -           Yes   LSP_
DEFAULT  3            xe3         No    172.4.5.5
```

P1

Following output displays the BGP routes of neighbor configured on P1.

```
P1#show ip bgp neighbors 172.1.2.53
BGP neighbor is 172.1.2.53, remote AS 100, local AS 100, internal link, peer index: 2
  BGP version 4, local router ID 200.200.200.54, remote router ID 44.44.44.53
  BGP state = Established, up for 01:14:22
  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 Labeled-Unicast: advertised and received
  Received 178 messages, 0 notifications, 0 in queue
  Sent 182 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 Labeled-Unicast  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Route-Reflector Client
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  1 accepted prefixes
  1 announced prefixes

Connections established 2; dropped 1
Local host: 172.1.2.54, Local port: 179
Foreign host: 172.1.2.53, Foreign port: 52756
TCP MSS: (0), Advertise TCP MSS: (1460), Send TCP MSS: (1460),  Receive TCP MSS: (536)
Sock FD : (26)
Nexthop: 172.1.2.54
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 01:14:22, due to Administratively Reset (Cease Notification sent)
Notification Error Message: (Cease/Administratively Reset.)
```

Following output displays the Labels attached to the BGP routes on P1.

```
P1#
P1#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 11.11.11.55/32   172.6.7.52        25600        25600
*>i 21.21.21.56/32   172.1.2.53        25601        25600


P1#show mpls forwarding-table | include 11.11.11.55

   B>   11.11.11.55/32      1         3         -           Yes   LSP_
DEFAULT  25600        xe21        No    172.6.7.52
```

Following output displays the Labels attached to the BGP routes on ABR2.

```
ABR2:
ABR2#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 11.11.11.55/32   172.1.2.54        25601        25600
*>i 21.21.21.56/32   172.3.4.56        25600        3


ABR2#show mpls forwarding-table | include 11.11.11.55

   B>   11.11.11.55/32      2         6         -           Yes   LSP_
DEFAULT  25600        xe6         No    172.1.2.54
```

PE2

Following output displays the PE2 configuration.

```
PE2#show running-config bgp
!
router bgp 100
 allocate-label all
 neighbor 11.11.11.55 remote-as 100
 neighbor 172.3.4.53 remote-as 100
```

```
 neighbor 11.11.11.55 update-source lo
 !
 address-family ipv4 unicast
 network 21.21.21.56/32
 exit-address-family
 !
 address-family ipv4 labeled-unicast
 bgp implicit-null
 neighbor 172.3.4.53 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 11.11.11.55 activate
 exit-address-family
 !
 address-family ipv4 vrf vrf1
 redistribute connected
 exit-address-family
 !
 exit
!
PE2#
```

Following output displays the Labels attached to the BGP routes on PE2.

```
PE2:
PE2#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 11.11.11.55/32   172.3.4.53        25600        25601
*>  21.21.21.56/32   0.0.0.0           3          -

PE2#show mpls forwarding-table | include 11.11.11.55
   B>   11.11.11.55/32      1         2         -           Yes   LSP_
DEFAULT  25601        ge10        No    172.3.4.53

PE2#show mpls ftn-table
 Primary FTN entry with FEC: 11.11.11.55/32, id: 1, row status: Active, Tunnel-Policy: N/A, State:
Installed
  Owner: BGP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  VRF id 0, BGP peer 172.3.4.53 BGP prefix 11.11.11.55
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
     Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 2 refcount: 1
      Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 2, owner: BGP, Stale: NO, refcount: 2, out intf: ge10, out label: 25601
    Nexthop addr: 172.3.4.53        cross connect ix: 1, op code: Push

PE2#show mpls vrf-table
Output for IPv4 VRF table with id: 2
 Primary FTN entry with FEC: 172.10.20.0/24, id: 1, row status: Active, Tunnel-Policy: N/A, State:
Installed
  Owner: BGP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  VRF id 2, BGP peer 11.11.11.55 BGP prefix 172.10.20.0
  Transport Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
     Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 1 refcount: 1
      Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: BGP, Stale: NO, refcount: 1, BGP out intf: ge10, transport out
intf: ge10, out label: 25601
    Nexthop addr: 11.11.11.55        cross connect ix: 2, op code: Push and Lookup


PE2#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP , K - CLI , V - LDP_VC, I - IGP_SHORTCUT
```

```
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

ILM-ECMP LDP: Disabled

Code    FEC/VRF/L2CKT    ILM-ID     In-Label    Out-Label   In-Intf    Out-
Intf/VRF       Nexthop          pri    LSP-Type
   B>   11.11.11.55/32     1          25600       25601        N/A         N/A                172.3.4
.53         Yes    LSP_DEFAULT
   B>   vrf1              2          25601       Nolabel      N/A         N/A                N/A
            Yes    LSP_DEFAULT
PE2#
```

Glossary

Key Terms/Acronym	Description
BGP	Border Gateway Protocol
LU	BGP Labeled-Unicast
MPLS	Multi Protocol Label Switching
PHP	Penultimate Hop Popping of MPLS label

© 2026 IP Infusion Inc. Proprietary

BGP MPLS Next Hop Tracking Configuration

MPLS nexthop tracking is used to track the registered next-hop in the NSM. It helps to notify the client when the LSP operational status changes, so that clients can take action accordingly.

BGP shall do lookup with NSM for LSP reachability, if lookup is successful, BGP marks the nexthop as reachable and register with NSM for the LSP tracking. Further proceeds for FTN/ILM installation and route advertisement. If it is down or does not exist, BGP will mark it as in-valid, route advertisement will not happen until LSP is reachable, BGP query NSM for every rib scan.

Notes:

- QOS should be enabled in all devices for this feature.
- Traceroute from CE-to-CE output doesn't show MPLS node details after enabling MPLS propagate-TTL in PE router. It only shows one hop between ingress to egress direction of PE routers.

Topology

Figure 81. BGP MPLS Next Hop Tracking

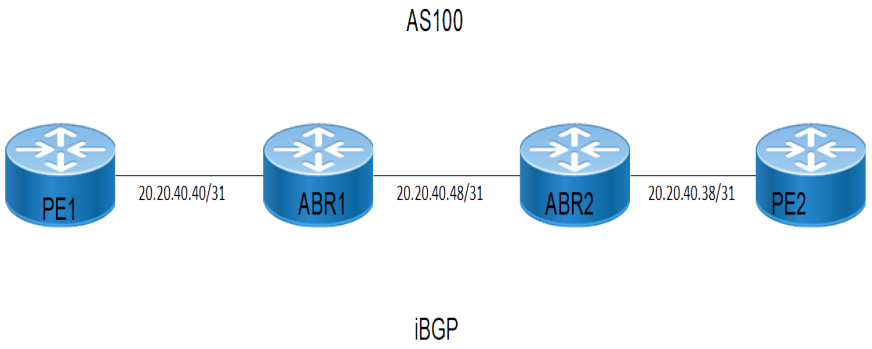

Configuration

PE1

`#configure terminal`	Enter the Configure mode.
`(config)#interface lo`	Enter interface mode
`(config-if)#ip address 1.1.1.1/32 secondary`	Configure the IP address of the interface loopback
`(config-if)#exit`	Exit interface mode
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)# router ldp`	Configure Router LDP instance
`(config)# commit`	Committing the configuration to apply in running configuration

© 2026 IP Infusion Inc. Proprietary

`(config)# bfd interval 3 minrx 3 multiplier 3`	Configure BFD single-hop sessions timer and reception interval in millisecond and the Hello multiplier.
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)#interface xe4`	Enter interface mode
`(config-if)#ip address 20.20.40.40/31`	Configure the IP address on the interface
`(config-if)#label-switching`	Enable label-switching on interface
`(config-if)# enable-ldp ipv4`	Enable LDP process on interface
`(config-if)#exit`	Exit interface mode.
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)#router ospf 1`	Configure the routing process OSPF with process id
`(config-router)# bfd all-interfaces`	Enable BFD for all neighbors.
`(config-router)#network 1.1.1.1/32 area 0`	Define the interface on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#network 20.20.40.40/31 area 0`	Define the interface on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#exit`	Exit from router ospf mode
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)#bgp nexthop-trigger enable`	Enable Nexthop tracking.
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)#router bgp 100`	Enter Router BGP mode
`(config-router)# mpls-nexthop-tracking`	Enable MPLS Next Hop tracking
`(config-router)#allocate-label all`	Allocate labels for all IPv4 prefixes advertised
`(config-router)#neighbor 4.4.4.4 remote-as 100`	Add loopback ip of PE2 as neighbor with neighbor AS
`(config-router)#neighbor 4.4.4.4 update- source lo`	Update the source for that particular neighbor as loopback interface
`(config-router)#neighbor 2.2.2.2 remote-as 100`	Add loopback ip of ABR1 as neighbor with neighbor AS
`(config-router)#neighbor 2.2.2.2 update- source lo`	Update the source for that particular neighbor as loopback interface
`(config-router)# neighbor 2.2.2.2 fall-over bfd multihop`	Enabling BFD at BGP level for particular neighbor.
`(config-router)# address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`(config-router-af)# network 1.1.1.1/32`	Advertise a route via iBGP connection.
`(config-router-af)#exit-address-family`	Exit from address family IPv4 unicast

© 2026 IP Infusion Inc. Proprietary

(config-router)#address-family ipv4 labeled-unicast	Enter into labeled-unicast address family
(config-router-af)#neighbor 2.2.2.2 activate	Activate the neighbor inside labeled-unicast address family
(config-router-af)#exit-address-family	Exit from address family IPv4 labeled unicast
(config-router)#address-family vpnv4 unicast	Enter into vpnv4 unicast address family
(config-router-af)#neighbor 4.4.4.4 activate	Activate the neighbor inside vpnv4 address family
(config-router-af)#exit-address-family	Exit from address family vpnv4
(config-router)#exit	Exit from router BGP mode
(config)# commit	Committing the configuration to apply in running configuration
(config)#end	Exit from config mode into privilege mode

ABR1

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 2.2.2.2/32 Secondary	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)# commit	Committing the configuration to apply in running configuration
(config)# router ldp	Configure Router LDP instance
(config)# commit	Committing the configuration to apply in running configuration
(config)# bfd interval 3 minrx 3 multiplier 3	Configure BFD single-hop sessions timer and reception interval in millisecond and the Hello multiplier.
(config)# commit	Committing the configuration to apply in running configuration
(config)#interface xe4	Enter interface mode
(config-if)#ip address 20.20.40.41/31	Configure the IP address on the interface
(config-if)#label-switching	Enable label-switching on interface
(config-if)# enable-ldp ipv4	Enable LDP process on interface
(config-if)#exit	Exit interface mode.
(config)#interface xe10	Enter interface mode
(config-if)#ip address 20.20.40.48/31	Configure the IP address on the interface
(config-if)#label-switching	Enable label-switching on interface
(config-if)# enable-ldp ipv4	Enable LDP process on interface
(config-if)#exit	Exit interface mode.
(config)# commit	Committing the configuration to apply in running

	configuration
`(config)# enable ext-ospf-multi-inst`	
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)#router ospf 1`	Configure the routing process OSPF with process id
`(config)# bfd all-interfaces`	Enable BFD for all neighbors.
`(config-router)#network 2.2.2.2/32 area 0`	Define the interface on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#network 20.20.40.40/31 area 0`	Define the interface on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#exit`	Exit from router ospf mode
`(config)#router ospf 2`	Configure the routing process OSPF with process id
`(config)# bfd all-interfaces`	Enable BFD for all neighbors.
`(config)# redistribute connected`	Redistribute connected routes.
`(config-router)#network 2.2.2.2/32 area 0 instance-id 2`	Advertise the network in Area 0 with an instance ID of 2.
`(config-router)#network 20.20.40.48/31 area 0`	Define the interface on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#exit`	Exit from router ospf mode
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)#bgp nexthop-trigger enable`	Enable Nexthop tracking.
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)#router bgp 100`	Enter Router BGP mode
`(config-router)# mpls-nexthop-tracking`	Enable MPLS Next Hop tracking
`(config-router)#allocate-label all`	Allocate labels for all IPv4 prefixes advertised
`(config-router)#neighbor 1.1.1.1 remote-as 100`	Add loopback ip of PE1 as neighbor with neighbor AS
`(config-router)#neighbor 1.1.1.1 update- source lo`	Update the source for that particular neighbor as loopback interface
`(config-router)# neighbor 1.1.1.1 fall-over bfd multihop`	Enabling BFD at BGP level for particular neighbor.
`(config-router)#neighbor 3.3.3.3 remote-as 100`	Add loopback ip of ABR2 as neighbor with neighbor AS
`(config-router)#neighbor 3.3.3.3 update- source lo`	Update the source for that particular neighbor as loopback interface
`(config-router)# neighbor 3.3.3.3 fall-over bfd multihop`	Enabling BFD at BGP level for particular neighbor.

© 2026 IP Infusion Inc. Proprietary

(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# network 2.2.2.2/32	Advertise a route via iBGP connection.
(config-router-af)#exit-address-family	Exit from address family IPv4 unicast
(config-router)#address-family ipv4 labeled-unicast	Enter into labeled-unicast address family
(config-router-af)#neighbor 1.1.1.1 activate	Activate the neighbor inside labeled-unicast address family
(config-router-af)# neighbor 1.1.1.1 route-reflector-client	Configure the router as Route-reflector (RR) and Neighbor router as its client.
(config-router-af)# neighbor 1.1.1.1 next-hop-self	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast.
(config-router-af)#neighbor 3.3.3.3 activate	Activate the neighbor inside labeled-unicast address family
(config-router-af)# neighbor 3.3.3.3 route-reflector-client	Configure the router as Route-reflector (RR) and Neighbor router as its client.
(config-router-af)# neighbor 3.3.3.3 next-hop-self	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast.
(config-router-af)#exit-address-family	Exit from address family IPv4 labeled unicast
(config-router)#exit	Exit from router BGP mode
(config)# commit	Committing the configuration to apply in running configuration
(config)#end	Exit from config mode into privilege mode

ABR2

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 3.3.3.3/32	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)# commit	Committing the configuration to apply in running configuration
(config)# router ldp	Configure Router LDP instance
(config)# commit	Committing the configuration to apply in running configuration
(config)# bfd interval 3 minrx 3 multiplier 3	Configure BFD single-hop sessions timer and reception interval in millisecond and the Hello multiplier.
(config)# commit	Committing the configuration to apply in running configuration
(config)#interface xe5	Enter interface mode
(config-if)#ip address 20.20.40.38/31	Configure the IP address on the interface

© 2026 IP Infusion Inc. Proprietary

(config-if)#label-switching	Enable label-switching on interface
(config-if)# enable-ldp ipv4	Enable LDP process on interface
(config-if)#exit	Exit interface mode.
(config)#interface xe10	Enter interface mode
(config-if)#ip address 20.20.40.49/31	Configure the IP address on the interface
(config-if)#label-switching	Enable label-switching on interface
(config-if)# enable-ldp ipv4	Enable LDP process on interface
(config-if)#exit	Exit interface mode.
(config)# commit	Committing the configuration to apply in running configuration
(config)# enable ext-ospf-multi-inst	
(config)# commit	Committing the configuration to apply in running configuration
(config)#router ospf 2	Configure the routing process OSPF with process id
(config)# bfd all-interfaces	Enable BFD for all neighbors.
(config)# redistribute connected	Redistribute connected routes.
(config-router)#network 3.3.3.3/32 area 0	Define the interface on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
(config-router)#network 20.20.40.48/31 area 0	Define the interface on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
(config-router)#exit	Exit from router ospf mode
(config)#router ospf 3	Configure the routing process OSPF with process id
(config)# bfd all-interfaces	Enable BFD for all neighbors.
(config-router)#network 2.2.2.2/32 area 0 instance-id 3	Advertise the network in Area 0 with an instance ID of 2.
(config-router)#network 20.20.40.38/31 area 0	Define the interface on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
(config-router)#exit	Exit from router ospf mode
(config)# commit	Committing the configuration to apply in running configuration
(config)#bgp nexthop-trigger enable	Enable Nexthop tracking.
(config)# commit	Committing the configuration to apply in running configuration
(config)#router bgp 100	Enter Router BGP mode
(config-router)# mpls-nexthop-tracking	Enable MPLS Next Hop tracking
(config-router)#allocate-label all	Allocate labels for all IPv4 prefixes advertised
(config-router)#neighbor 4.4.4.4 remote-as 100	Add loopback ip of PE2 as neighbor with neighbor AS

© 2026 IP Infusion Inc. Proprietary

(config-router)#neighbor 4.4.4.4 update- source lo	Update the source for that particular neighbor as loopback interface
(config-router)# neighbor 4.4.4.4 fall-over bfd multihop	Enabling BFD at BGP level for particular neighbor.
(config-router)#neighbor 2.2.2.2 remote-as 100	Add loopback ip of ABR1 as neighbor with neighbor AS
(config-router)#neighbor 2.2.2.2 update- source lo	Update the source for that particular neighbor as loopback interface
(config-router)# neighbor 2.2.2.2 fall-over bfd multihop	Enabling BFD at BGP level for particular neighbor.
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# network 2.2.2.2/32	Advertise a route via iBGP connection.
(config-router-af)#exit-address-family	Exit from address family IPv4 unicast
(config-router)#address-family ipv4 labeled-unicast	Enter into labeled-unicast address family
(config-router-af)#neighbor 2.2.2.2 activate	Activate the neighbor inside labeled-unicast address family
(config-router-af)# neighbor 2.2.2.2 route-reflector-client	Configure the router as Route-reflector (RR) and Neighbor router as its client.
(config-router-af)# neighbor 2.2.2.2 next-hop-self	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast.
(config-router-af)#neighbor 4.4.4.4 activate	Activate the neighbor inside labeled-unicast address family
(config-router-af)# neighbor 4.4.4.4 route-reflector-client	Configure the router as Route-reflector (RR) and Neighbor router as its client.
(config-router-af)# neighbor 4.4.4.4 next-hop-self	Enable next hop self for the particular neighbor inside address family IPv4 labeled unicast.
(config-router-af)#exit-address-family	Exit from address family IPv4 labeled unicast
(config-router)#exit	Exit from router BGP mode
(config)# commit	Committing the configuration to apply in running configuration
(config)#end	Exit from config mode into privilege mode

PE2

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 4.4.4.4/32 secondary	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)# commit	Committing the configuration to apply in running configuration

© 2026 IP Infusion Inc. Proprietary

`(config)# router ldp`	Configure Router LDP instance
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)# bfd interval 3 minrx 3 multiplier 3`	Configure BFD single-hop sessions timer and reception interval in millisecond and the Hello multiplier.
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)#interface xe5`	Enter interface mode
`(config-if)#ip address 20.20.40.39/31`	Configure the IP address on the interface
`(config-if)#label-switching`	Enable label-switching on interface
`(config-if)# enable-ldp ipv4`	Enable LDP process on interface
`(config-if)#exit`	Exit interface mode.
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)#router ospf 3`	Configure the routing process OSPF with process id
`(config)# bfd all-interfaces`	Enable BFD for all neighbors.
`(config-router)#network 3.3.3.3/32 area 0`	Define the interface on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#network 20.20.40.38/31 area 0`	Define the interface on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#exit`	Exit from router ospf mode
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)#bgp nexthop-trigger enable`	Enable Nexthop tracking.
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)#router bgp 100`	Enter Router BGP mode
`(config-router)# mpls-nexthop-tracking`	Enable MPLS Next Hop tracking
`(config-router)#allocate-label all`	Allocate labels for all IPv4 prefixes advertised
`(config-router)#neighbor 1.1.1.1 remote-as 100`	Add loopback ip of PE1 as neighbor with neighbor AS
`(config-router)#neighbor 1.1.1.1 update- source lo`	Update the source for that particular neighbor as loopback interface
`(config-router)# neighbor 1.1.1.1 fall-over bfd multihop`	Enabling BFD at BGP level for particular neighbor.
`(config-router)#neighbor 3.3.3.3 remote-as 100`	Add loopback ip of ABR2 as neighbor with neighbor AS
`(config-router)#neighbor 3.3.3.3 update- source lo`	Update the source for that particular neighbor as loopback interface

© 2026 IP Infusion Inc. Proprietary

(config-router)# neighbor 3.3.3.3 fall-over bfd multihop	Enabling BFD at BGP level for particular neighbor.
(config-router)# address-family ipv4 unicast	Enter address-family ipv4 unicast mode
(config-router-af)# network 4.4.4.4/32	Advertise a route via iBGP connection.
(config-router-af)#exit-address-family	Exit from address family IPv4 unicast
(config-router)#address-family ipv4 labeled-unicast	Enter into labeled-unicast address family
(config-router-af)#neighbor 3.3.3.3 activate	Activate the neighbor inside labeled-unicast address family
(config-router-af)#exit-address-family	Exit from address family IPv4 labeled unicast
(config-router)#address-family vpnv4 unicast	Enter into vpnv4 unicast address family
(config-router-af)#neighbor 1.1.1.1 activate	Activate the neighbor inside vpnv4 address family
(config-router-af)#exit-address-family	Exit from address family vpnv4
(config-router)#exit	Exit from router BGP mode
(config)# commit	Committing the configuration to apply in running configuration
(config)#end	Exit from config mode into privilege mode

Validation

PE1

```
PE1#show ip bgp labeled-unicast summary
BGP router identifier 1.1.1.1, 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
2.2.2.2                  4   100   90         89       3      0      0  00:36:22               3

Total number of neighbors 1

Total number of Established sessions 1


PE1#show bgp nexthop-tracking
Configured NHT: ENABLED
NHT Delay time-interval : 5
Configured MPLS NHT: ENABLED
BGP VRF: (Default) VRF_ID 0
BGP Instance: (Default), AS: 100, router-id 1.1.1.1
NHT is Enabled
Recvd Msg count from RIB: 0
NHT delay-timer remaining seconds: 0
BGP nexthop(s):
 2.2.2.2        Registered (MPLS also Registered)
Total number of IPV4 nexthops : 1
Total number of IPV6 nexthops : 0

PE1#show ldp session
Peer IP Address           IF Name    My Role    State      KeepAlive  UpTime
2.2.2.2                   xe4        Passive   OPERATIONAL   30    00:38:06
```

```
PE1#show ip bgp vpnv4 all summary
BGP router identifier 1.1.1.1, 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
4.4.4.4                  4   100   92         93       1      0      0  00:38:49               0

Total number of neighbors 1

Total number of Established sessions 1

PE1#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      00:00:34    20.20.40.41     xe4                      0

PE1#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, 00:59:57
O            2.2.2.2/32 [110/2] via 20.20.40.41, xe4, 00:40:05
B            3.3.3.3/32 [200/0] via 2.2.2.2 (recursive via 20.20.40.41), 00:36:31
B            4.4.4.4/32 [200/0] via 2.2.2.2 (recursive via 20.20.40.41), 00:36:31
C            20.20.40.40/31 is directly connected, xe4, 00:58:24
C            127.0.0.0/8 is directly connected, lo, 01:06:14

Gateway of last resort is not set
```

PE2

```
PE2#show ip bgp labeled-unicast summary
BGP router identifier 4.4.4.4, 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/Dow
n   State/PfxRcd
3.3.3.3                  4   100  103        102       3      0      0  00:41:58
              3

Total number of neighbors 1

Total number of Established sessions 1
PE2#show bgp nexthop-tracking
Configured NHT: ENABLED
NHT Delay time-interval : 5
Configured MPLS NHT: ENABLED
BGP VRF: (Default) VRF_ID 0
BGP Instance: (Default), AS: 100, router-id 4.4.4.4
NHT is Enabled
Recvd Msg count from RIB: 0
NHT delay-timer remaining seconds: 0
```

```
BGP nexthop(s):
 3.3.3.3        Registered (MPLS also Registered)
Total number of IPV4 nexthops : 1
Total number of IPV6 nexthops : 0

PE2#show ldp session
Peer IP Address           IF Name    My Role    State      KeepAlive  UpTime
3.3.3.3                   xe5        Active    OPERATIONAL   30    00:42:33
PE2#show bgp vpn
PE2#show ip bgp vpnv4 all summary
BGP router identifier 4.4.4.4, 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/Dow
n   State/PfxRcd
1.1.1.1                  4   100   99         99       1      0      0  00:41:48
               0

Total number of neighbors 1

Total number of Established sessions 1
PE2#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 3 VRF(default):
Neighbor ID     Pri   State            Dead Time   Address         Interface
       Instance ID
3.3.3.3           1   Full/Backup      00:00:35    20.20.40.38     xe5
           0
PE2#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 [200/0] via 3.3.3.3 (recursive via 20.20.40.38), 00:42:1
7
B            2.2.2.2/32 [200/0] via 3.3.3.3 (recursive via 20.20.40.38), 00:42:1
7
O            3.3.3.3/32 [110/2] via 20.20.40.38, xe5, 00:44:10
C            4.4.4.4/32 is directly connected, lo, 01:05:50
C            20.20.40.38/31 is directly connected, xe5, 01:01:12
C            127.0.0.0/8 is directly connected, lo, 01:11:16

Gateway of last resort is not set
```

ABR1

```
ABR1#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          00:00:34    20.20.40.40     xe4                      0

Total number of full neighbors: 1
OSPF process 2 VRF(default):
Neighbor ID     Pri   State            Dead Time   Address         Interface           Instance ID
```

```
3.3.3.3           1   Full/Backup      00:00:37    20.20.40.49     xe10                    0
ABR1#show ldp session
Peer IP Address           IF Name    My Role    State      KeepAlive  UpTime
1.1.1.1                   xe4        Active    OPERATIONAL   30    00:46:10
3.3.3.3                   xe10       Passive   OPERATIONAL   30    00:43:55
ABR1#sh ip bgp vpnv4 all summary
ABR1#sh ip bgp labeled-unicast summary
BGP router identifier 2.2.2.2, 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
1.1.1.1                  4   100  113        113       2      0      0  00:46:44               1
3.3.3.3                  4   100  109        111       2      0      0  00:45:10               2

Total number of neighbors 2

Total number of Established sessions 2
ABR1#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            1.1.1.1/32 [110/2] via 20.20.40.40, xe4, 00:47:19
C            2.2.2.2/32 is directly connected, lo, 01:07:16
O            3.3.3.3/32 [110/2] via 20.20.40.49, xe10, 00:45:26
B            4.4.4.4/32 [200/0] via 3.3.3.3 (recursive via 20.20.40.49), 00:43:51
O E2         20.20.40.38/31 [110/20] via 20.20.40.49, xe10, 00:45:26
C            20.20.40.40/31 is directly connected, xe4, 00:58:36
C            20.20.40.48/31 is directly connected, xe10, 00:58:23
C            127.0.0.0/8 is directly connected, lo, 01:14:00

Gateway of last resort is not set

ABR1#show bgp nexthop-tracking
Configured NHT: ENABLED
NHT Delay time-interval : 5
Configured MPLS NHT: ENABLED
BGP VRF: (Default) VRF_ID 0
BGP Instance: (Default), AS: 100, router-id 2.2.2.2
NHT is Enabled
Recvd Msg count from RIB: 0
NHT delay-timer remaining seconds: 0
BGP nexthop(s):
 1.1.1.1        Registered (MPLS also Registered)
 3.3.3.3        Registered (MPLS also Registered)
Total number of IPV4 nexthops : 2
Total number of IPV6 nexthops : 0
```

ABR2

```
ABR2#show 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/DR         00:00:37    20.20.40.48     xe10                    0
```

```
Total number of full neighbors: 1
OSPF process 3 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface           Instance ID
4.4.4.4           1   Full/DR         00:00:36    20.20.40.39     xe5                     0
ABR2#
ABR2#show ldp session
Peer IP Address           IF Name    My Role    State      KeepAlive  UpTime
2.2.2.2                   xe10       Active    OPERATIONAL   30    00:46:31
4.4.4.4                   xe5        Passive   OPERATIONAL   30    00:47:20
ABR2#show ip bgp labeled-unicast summary
BGP router identifier 3.3.3.3, 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
2.2.2.2                  4   100  115        114       2      0      0  00:47:21               2
4.4.4.4                  4   100  114        115       2      0      0  00:47:21               1

Total number of neighbors 2

Total number of Established sessions 2
ABR2#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 [200/0] via 2.2.2.2 (recursive via 20.20.40.48), 00:46:06
O            2.2.2.2/32 [110/2] via 20.20.40.48, xe10, 00:47:51
C            3.3.3.3/32 is directly connected, lo, 01:09:37
O            4.4.4.4/32 [110/2] via 20.20.40.39, xe5, 00:47:51
C            20.20.40.38/31 is directly connected, xe5, 00:48:33
O E2         20.20.40.40/31 [110/20] via 20.20.40.48, xe10, 00:47:51
C            20.20.40.48/31 is directly connected, xe10, 00:48:44
C            127.0.0.0/8 is directly connected, lo, 01:15:15

Gateway of last resort is not set
ABR2#show bgp nexthop-tracking
Configured NHT: ENABLED
NHT Delay time-interval : 5
Configured MPLS NHT: ENABLED
BGP VRF: (Default) VRF_ID 0
BGP Instance: (Default), AS: 100, router-id 3.3.3.3
NHT is Enabled
Recvd Msg count from RIB: 0
NHT delay-timer remaining seconds: 0
BGP nexthop(s):
 2.2.2.2        Registered (MPLS also Registered)
 4.4.4.4        Registered (MPLS also Registered)
Total number of IPV4 nexthops : 2
Total number of IPV6 nexthops : 0
```

© 2026 IP Infusion Inc. Proprietary

BGP Labeled Unicast Next Hop Self in Route-Map

Overview

The BGP-LU next-hop-self in route map feature provides the ability set the local BGP peer as the next-hop for select BGP-LU routes. When such a route map is applied to a BGP-LU neighbor in the outbound direction, the matched routes will be updated as below:

- The next-hop address is replaced with the local BGP peer address. Based on the BGP configuration, this is either the local interface address or the local loopback address
- The label is replaced with the local label allocated for the prefix

Feature Characteristics

BGP-LU routes permitted by the route-map and configured with `set ip next-hop self` are advertised with the local BGP peer address as the next-hop and the locally assigned label. BGP-LU routes permitted by the route-map without `set next-hop self` are advertised with their original next-hop and the label remains unchanged. The BGP-LU routes denied by the route-map are not advertised.

Limitations

The existing route-map CLIs to set next-hop, such as, `set ip next-hop a.b.c.d` or `set ip next-hop peer-address` are not recommended to be used for this feature and they can impact adversely.

Prerequisites

While deploying BGP labelled unicast, ensure that the `allocate-label` command is enabled under `router bgp` mode. This command is required for allocating labels to IPv4 prefixes.

Configuration

Topology

This topology contains Route Reflectors (RR) and PE nodes with BGP-LU as the transport between them.

Domain 1 contains LDP as transport with OSPF as IGP on all nodes. Domain 2 contains RSVP as transport with IS-IS as IGP on all nodes.

© 2026 IP Infusion Inc. Proprietary

Figure 82. BGP LU Next-hop in Route-map

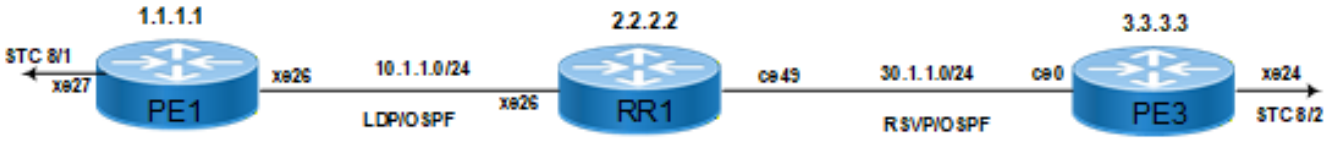

To configure BGP LU Next-hop, follow the steps mentioned below:

1. Configure BGP LU as transport

```
#configure terminal
(config)#interface lo
(config-if)#ip address 11.11.11.55/32 secondary
(config-if)#exit
(config)#interface xe16
(config-if)#ip address 172.4.5.55/24
(config-if)#label-switching
(config-if)#exit
(config)#commit
```

2. Configure the routing process OSPF with process ID 1.

```
(config)#router ospf 1
```

3. Define the interface (172.4.5.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).

```
 (config-router)#network 172.4.5.0/24 area 0
```

4. Define the interface (11.11.11.55/32) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).

```
 (config-router)#network 11.11.11.55/32 area 0
(config-router)#exit
(config)#commit
(config)#router bgp 100
```

5. Create a Prefix-list to match the advertised prefixes

```
 (config-router)# ip prefix-list BGP-LU-FILTER permit 10.10.10.0/24
```

6. Create Route-map for LU using the prefix-list

```
 (config-router)#route-map RM-BGPLU-OUT permit 10
(config-router)#match ip address prefix-list BGP-LU-FILTER
```

Optionally, apply strict filtering using `deny`

```
 (config-router)#route-map RM-BGPLU-OUT deny 20
```

7. Apply Route-map on BGP-LU Neighbor

```
(config-router)# allocate-label all route-map RM-BGPLU-OUT
```

Configuration Snapshot

PE1

```
PE1#sh run
!
service password-encryption
!
logging console 3
logging monitor 5
logging logfile device_debug_log 2
logging level nsm 5
logging level ospf 5
logging level ldp 5
logging level hsl 5
logging level bgp 5
logging level cml 5
logging level cmm 4
logging level all 5
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
bgp extended-asn-cap
!
forwarding profile kaps profile-two
hardware-profile filter qos enable
hardware-profile statistics ingress-acl enable
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
hostname PE1
tfo Disable
errdisable cause stp-bpdu-guard
feature dns relay
ip dns relay
ipv6 dns relay
feature rsyslog
logging remote server 10.16.58.70 5 port 1514 vrf management
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!
ip vrf management
!
ip vrf VRF1
 rd 100:1
 route-target both 100:1
!
ip vrf VRF2
```

```
 rd 100:2
 route-target both 100:2
!
router ldp
 router-id 1.1.1.1
 session-protection
 pw-status-tlv
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface ge1
!
interface ge2
!
interface ge3
!
interface ge4
!
interface ge5
!
interface ge6
!
interface ge7
!
interface ge8
!
interface ge9
!
interface ge10
!
interface ge11
!
interface ge12
!
interface ge13
!
interface ge14
!
interface ge15
!
interface ge16
!
interface ge17
!
interface ge18
!
interface ge19
!
interface ge20
!
interface ge21
!
interface ge22
!
interface ge29
!
interface lo
 ip address 127.0.0.1/8
 ip address 1.1.1.1/32 secondary
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
```

© 2026 IP Infusion Inc. Proprietary

```
interface xe23
!
interface xe24
!
interface xe25
!
interface xe26
 load-interval 30
 ip address 10.1.1.1/24
 mtu 9216
 label-switching
 ip ospf network point-to-point
 enable-ldp ipv4
!
interface xe27
 load-interval 30
 mtu 9216
!
interface xe27.102
 description L3VPN-VRF1
 encapsulation dot1q 102
 load-interval 30
 ip vrf forwarding VRF1
 ip address 101.1.1.1/24
 isis network point-to-point
 ip router isis 100
!
interface xe27.103
 description L3VPN-VRF2
 encapsulation dot1q 103
 load-interval 30
 ip vrf forwarding VRF2
 ip address 101.1.2.1/24
 isis network point-to-point
 ip router isis 200
!
interface xe28
!
 exit
!
router ospf 65535
 ospf router-id 1.1.1.1
 bfd all-interfaces
 network 1.1.1.1/32 area 0.0.0.0
 network 10.1.1.0/24 area 0.0.0.0
!
router isis 100 VRF1
 is-type level-1-2
 metric-style wide
 dynamic-hostname
 bfd all-interfaces
 net 49.0001.0000.0000.0001.00
 redistribute bgp
!
router isis 200 VRF2
 is-type level-1-2
 metric-style wide
 dynamic-hostname
 bfd all-interfaces
 net 49.0002.0000.0000.0002.00
 redistribute bgp
!
router bgp 4200000001
 bgp router-id 1.1.1.1
 bgp auto-policy-soft-reset enable
 bgp log-neighbor-changes
 no bgp inbound-route-filter
 allocate-label all
```

© 2026 IP Infusion Inc. Proprietary

```
 neighbor 2.2.2.2 remote-as 4200000001
 neighbor 2.2.2.2 tcp-mss 1440
 neighbor 2.2.2.2 update-source 1.1.1.1
 neighbor 2.2.2.2 authentication-key
0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157be7f6e35f32bf637
 neighbor 2.2.2.2 advertisement-interval 0
 neighbor 2.2.2.2 fall-over bfd multihop
 !
 address-family ipv4 unicast
 network 1.1.1.1/32
 exit-address-family
 !
 address-family ipv4 labeled-unicast
 neighbor 2.2.2.2 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 2.2.2.2 activate
 exit-address-family
 !
 address-family ipv4 vrf VRF1
 redistribute connected
 redistribute isis
 exit-address-family
 !
 address-family ipv4 vrf VRF2
 redistribute connected
 redistribute isis
 exit-address-family
 !
 exit
!
line console 0
 exec-timeout 0 0
line vty 0 16
 exec-timeout 0 0
!
!
end
```

PE2

```
PE2-7001#sh running-config
!
! Software version: UFI_S9600-56DX-OcNOS-SP-PLUS-7.0.0.168-Alpha
11/13/202                                                                    5 18:32:31
!
! Last configuration change at 11:31:56 UTC Tue Nov 18 2025 by root
!
feature netconf-ssh
feature netconf-tls
!
feature netconf notification-cache enable
max-cache-notifications 0
!
cml bulk-config limit cpu enable
!
background-debug
 log all
  level 7
  suppress-non-bdr-logs
  exit
 !
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
snmp-server enable traps pwdelete
snmp-server enable traps pw
```

```
snmp-server enable traps mpls
snmp-server enable traps snmp authentication
snmp-server enable traps ospf
snmp-server enable traps bgp
snmp-server enable traps isis
!
bgp extended-asn-cap
!
hardware-profile filter evpn-mpls-mh enable
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
hardware-profile port-config mode3
!
bfd interval 3 minrx 3 multiplier 3
!
key chain isis
 key-id 3
  key-string encrypted 0xc8a471564ac751dc
!
key chain BGP
 key-id 4
  key-string encrypted
0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157b
                            e7f6e35f32bf637
!
qos enable
qos statistics
!
mpls ilm-ecmp ldp
mpls ftn-ecmp ldp
mpls label mode vpnv4 all-vrfs per-vrf
mpls label mode vpnv6 all-vrfs per-vrf
!
mpls vpls vpls4294961250 1250
 control-word
 signaling bgp
  ve-id 61250
 exit-signaling
 exit-vpls
!
mpls vpls vpls4294961251 1251
 control-word
 signaling bgp
  ve-id 61251
 exit-signaling
 exit-vpls
!
mpls vpls vpls4294961252 1252
 control-word
 signaling bgp
  ve-id 61252
 exit-signaling
 exit-vpls
!
mpls vpls vpls4294961253 1253
 control-word
 signaling bgp
  ve-id 61253
 exit-signaling
 exit-vpls
!
mpls vpls vpls4294961254 1254
 control-word
 signaling bgp
  ve-id 61254
 exit-signaling
 exit-vpls
!
```

```
mpls vpls vpls4294961255 1255
 control-word
 signaling bgp
  ve-id 61255
 exit-signaling
 exit-vpls
!
mpls vpls vpls4294961256 1256
 control-word
 signaling bgp
  ve-id 61256
 exit-signaling
 exit-vpls
!
mpls vpls vpls4294961257 1257
 control-word
 signaling bgp
  ve-id 61257
 exit-signaling
 exit-vpls
!
mpls vpls vpls4294961258 1258
 control-word
 signaling bgp
  ve-id 61258
 exit-signaling
 exit-vpls
!
mpls vpls vpls4294961259 1259
 control-word
 signaling bgp
  ve-id 61259
 exit-signaling
 exit-vpls
!
mpls vpls vpls4294961260 1260
 control-word
 signaling bgp
  ve-id 61260
 exit-signaling
 exit-vpls
!
mpls vpls VPLS_AD1301 61301
 vpls-mtu 8000
 control-word
 signaling ldp
  vpls-type ethernet
  bgp-auto-discovery
    l2vpn-id 192.168.36.2:1301
    rd 4294967294:61301
    route-target both 4294967294:61301
  exit-bgp-auto-discovery
 exit-signaling
 exit-vpls
!
mpls vpls vpls1261 4294961261
 control-word
 signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
 exit-signaling
 exit-vpls
!
mpls vpls vpls1262 4294961262
 control-word
 signaling ldp
  vpls-peer 192.168.36.1
```

```
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
 exit-signaling
 exit-vpls
!
mpls vpls vpls1263 4294961263
 control-word
 signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
 exit-signaling
 exit-vpls
!
mpls vpls vpls1264 4294961264
 control-word
 signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
 exit-signaling
 exit-vpls
!
mpls vpls vpls1265 4294961265
 control-word
 signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
 exit-signaling
 exit-vpls
!
mpls vpls vpls1266 4294961266
 control-word
 signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
 exit-signaling
 exit-vpls
!
mpls vpls vpls1267 4294961267
 control-word
 signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
 exit-signaling
 exit-vpls
!
mpls vpls vpls1268 4294961268
 control-word
 signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
 exit-signaling
 exit-vpls
!
mpls vpls vpls1269 4294961269
 control-word
 signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
 exit-signaling
 exit-vpls
!
```

```
mpls vpls vpls1270 4294961270
 control-word
 signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
 exit-signaling
 exit-vpls
!
mpls vpls VPLS_AD1302 4294961302
 vpls-mtu 8000
 control-word
 signaling ldp
  vpls-type ethernet
  bgp-auto-discovery
    l2vpn-id 65000:4294961302
    rd 4294967294:61302
    route-target both 4294967294:61302
  exit-bgp-auto-discovery
 exit-signaling
 exit-vpls
!
mpls vpls VPLS_AD1303 4294961303
 vpls-mtu 8000
 control-word
 signaling ldp
  vpls-type ethernet
  bgp-auto-discovery
    l2vpn-id 65000:4294961303
    rd 4294967294:61303
    route-target both 4294967294:61303
  exit-bgp-auto-discovery
 exit-signaling
 exit-vpls
!
mpls vpls VPLS_AD1304 4294961304
 vpls-mtu 8000
 control-word
 signaling ldp
  vpls-type ethernet
  bgp-auto-discovery
    l2vpn-id 65000:4294961304
    rd 4294967294:61304
    route-target both 4294967294:61304
  exit-bgp-auto-discovery
 exit-signaling
 exit-vpls
!
mpls vpls VPLS_AD1305 4294961305
 vpls-mtu 8000
 control-word
 signaling ldp
  vpls-type ethernet
  bgp-auto-discovery
    l2vpn-id 65000:4294961305
    rd 4294967294:61305
    route-target both 4294967294:61305
  exit-bgp-auto-discovery
 exit-signaling
 exit-vpls
!
mpls vpls VPLS_AD1306 4294961306
 vpls-mtu 8000
 control-word
 signaling ldp
  vpls-type ethernet
  bgp-auto-discovery
    l2vpn-id 65000:4294961306
```

```
    rd 4294967294:61306
    route-target both 4294967294:61306
  exit-bgp-auto-discovery
 exit-signaling
 exit-vpls
!
mpls vpls VPLS_AD1307 4294961307
 vpls-mtu 8000
 control-word
 signaling ldp
  vpls-type ethernet
  bgp-auto-discovery
    l2vpn-id 65000:4294961307
    rd 4294967294:61307
    route-target both 4294967294:61307
  exit-bgp-auto-discovery
 exit-signaling
 exit-vpls
!
mpls vpls VPLS_AD1308 4294961308
 vpls-mtu 8000
 control-word
 signaling ldp
  vpls-type ethernet
  bgp-auto-discovery
    l2vpn-id 65000:4294961308
    rd 4294967294:61308
    route-target both 4294967294:61308
  exit-bgp-auto-discovery
 exit-signaling
 exit-vpls
!
mpls vpls VPLS_AD1309 4294961309
 vpls-mtu 8000
 control-word
 signaling ldp
  vpls-type ethernet
  bgp-auto-discovery
    l2vpn-id 65000:4294961309
    rd 4294967294:61309
    route-target both 4294967294:61309
  exit-bgp-auto-discovery
 exit-signaling
 exit-vpls
!
mpls vpls VPLS_AD1310 4294961310
 vpls-mtu 8000
 control-word
 signaling ldp
  vpls-type ethernet
  bgp-auto-discovery
    l2vpn-id 65000:4294961310
    rd 4294967294:61310
    route-target both 4294967294:61310
  exit-bgp-auto-discovery
 exit-signaling
 exit-vpls
!
mpls l2-circuit PE2-To-PE1-1271 1271 192.168.36.1
 control-word
!
mpls l2-circuit PE2-To-PE1-1272 4294961272 192.168.36.1
 control-word
!
mpls l2-circuit PE2-To-PE1-1273 4294961273 192.168.36.1
 control-word
!
mpls l2-circuit PE2-To-PE1-1274 4294961274 192.168.36.1
```

```
 control-word
!
mpls l2-circuit PE2-To-PE1-1275 4294961275 192.168.36.1
 control-word
!
mpls l2-circuit PE2-To-PE1-1276 4294961276 192.168.36.1
 control-word
!
mpls l2-circuit PE2-To-PE1-1277 4294961277 192.168.36.1
 control-word
!
mpls l2-circuit PE2-To-PE1-1278 4294961278 192.168.36.1
 control-word
!
mpls l2-circuit PE2-To-PE1-1279 4294961279 192.168.36.1
 control-word
!
mpls l2-circuit PE2-To-PE1-1280 4294961280 192.168.36.1
 control-word
!
mpls l2-circuit PE2-To-PE3-1281 4294961281 192.168.36.3
 control-word
!
mpls l2-circuit PE2-To-PE3-1282 4294961282 192.168.36.3
 control-word
!
mpls l2-circuit PE2-To-PE3-1283 4294961283 192.168.36.3
 control-word
!
mpls l2-circuit PE2-To-PE3-1284 4294961284 192.168.36.3
 control-word
!
mpls l2-circuit PE2-To-PE3-1285 4294961285 192.168.36.3
 control-word
!
mpls l2-circuit PE2-To-PE3-1286 4294961286 192.168.36.3
 control-word
!
mpls l2-circuit PE2-To-PE3-1287 4294961287 192.168.36.3
 control-word
!
mpls l2-circuit PE2-To-PE3-1288 4294961288 192.168.36.3
 control-word
!
mpls l2-circuit PE2-To-PE3-1289 4294961289 192.168.36.3
 control-word
!
mpls l2-circuit PE2-To-PE3-1290 4294961290 192.168.36.3
 control-word
!
mpls l2-circuit PE2-To-PE4-1291 4294961291 192.168.36.4
 control-word
!
mpls l2-circuit PE2-To-PE4-1292 4294961292 192.168.36.4
 control-word
!
mpls l2-circuit PE2-To-PE4-1293 4294961293 192.168.36.4
 control-word
!
mpls l2-circuit PE2-To-PE4-1294 4294961294 192.168.36.4
 control-word
!
mpls l2-circuit PE2-To-PE4-1295 4294961295 192.168.36.4
 control-word
!
mpls l2-circuit PE2-To-PE4-1296 4294961296 192.168.36.4
 control-word
!
```

© 2026 IP Infusion Inc. Proprietary

```
mpls l2-circuit PE2-To-PE4-1297 4294961297 192.168.36.4
 control-word
!
mpls l2-circuit PE2-To-PE4-1298 4294961298 192.168.36.4
 control-word
!
mpls l2-circuit PE2-To-PE4-1299 4294961299 192.168.36.4
 control-word
!
mpls l2-circuit PE2-To-PE4-1300 4294961300 192.168.36.4
 control-word
!
mpls l2-circuit vc2000 4294967290 192.168.36.11
!
mpls l2-circuit vc2001 4294967290 192.168.36.12
!
mpls l2-circuit vc2004 4294967291 192.168.36.11
!
mpls l2-circuit vc2003 4294967291 192.168.36.12
!
hostname PE2-7001
port ce46 breakout 4X10g
port ce47 breakout 4X10g
ip name-server vrf management 10.12.3.23
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
ospf restart grace-period 2
ospf restart helper max-grace-period 2
aaa local authentication attempts max-fail 25
aaa local authentication unlock-timeout 1
aaa authentication login error-enable
snmp-server enable snmp vrf management
snmp-server view all .1 included vrf management
snmp-server community public vrf management
feature dns relay
ip dns relay
ipv6 dns relay
feature rsyslog
logging remote server 10.16.100.20 5 port 1514 vrf management
logging remote server 10.16.100.20 5 port 1514
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!
router-id 192.168.36.2
!
evpn mpls enable
!
evpn mpls irb
!
ip vrf irbvrf3001
 rd 4294967294:63001
 route-target both 4294967294:63001
 l3vni 23001
 maximum-fib-routes ipv4 10000 warning-only
 maximum-fib-routes ipv6 10000 warning-only
!
ip vrf irbvrf3002
 rd 192.168.36.2:3002
 route-target import 4294967294:63002
 route-target export 4294967294:63022
 l3vni 23002
```

```
!
ip vrf irbvrf3003
 rd 192.168.36.2:3003
 route-target import 4294967294:63003
 route-target export 4294967294:63033
 export map RM-EXPORT-EVPN-IRBVRF3003-IPv6
 l3vni 23003
!
ip vrf management
!
ip vrf vrf101
 rd 4294967294:101
 route-target both 4294967294:101
 maximum-fib-routes ipv4 10000 stop-install
 maximum-fib-routes ipv6 10000 stop-install
!
ip vrf vrf102
 rd 4294967294:100
 route-target both 4294967294:102
 export map RM-EXPORT-ATTR
!
ip vrf vrf103
 rd 4294967294:103
 route-target both 4294967294:103
 export map RM-EXPORT-EVPN-IRBVRF3003-IPv6
!
ip vrf vrf104
 rd 4294967294:104
 route-target both 4294967294:104
!
ip vrf vrf105
 rd 4294967294:105
 route-target both 4294967294:105
!
ip vrf vrf106
 rd 4294967294:106
 route-target both 4294967294:106
!
ip vrf vrf107
 rd 4294967294:107
 route-target both 4294967294:107
!
ip vrf vrf108
 rd 4294967294:108
 route-target both 4294967294:108
!
ip vrf vrf109
 rd 4294967294:109
 route-target both 4294967294:109
!
ip vrf vrf110
 rd 4294967294:11
 route-target both 4294967294:110
 maximum-fib-routes ipv4 10000 stop-install
 maximum-fib-routes ipv6 10000 stop-install
!
ip vrf vrf111
 rd 4294967294:111
 route-target both 4294967294:111
 export map RM-EXPORT-VRF111
!
mac vrf ELAN1901
 rd 192.168.36.2:1901
 route-target both 1901:1901
 export map RM-EXPORT-ELAN1901
!
mac vrf ELAN1902
 rd 4294967294:61902
```

```
 route-target both 4294967294:61902
!
mac vrf ELAN1903
 rd 4294967294:1903
 route-target both 4294967294:1903
!
mac vrf ELAN1904
 rd 4294967294:61904
 route-target both 4294967294:41904
!
mac vrf ELAN1905
 rd 4294967294:1905
 route-target both 4294967294:21901
!
mac vrf ELAN1906
 rd 4294967294:1906
 route-target both 4294967294:41906
!
mac vrf ELAN1907
 rd 4294967294:1907
 route-target both 4294967294:41907
!
mac vrf ELAN1908
 rd 4294967294:1908
 route-target both 4294967294:41908
!
mac vrf ELAN1909
 rd 4294967294:1909
 route-target both 4294967294:41909
!
mac vrf ELAN1910
 rd 4294967294:1910
 route-target both 4294967294:41910
!
mac vrf eline901
 rd 192.168.36.2:901
 route-target both 901:901
!
mac vrf eline902
 rd 4294967294:902
 route-target both 4294967294:902
!
mac vrf eline903
 rd 4294967294:903
 route-target both evpn-auto-rt
!
mac vrf eline904
 rd 192.168.36.2:904
 route-target both evpn-auto-rt
 route-target both 904:904
!
mac vrf eline905
 rd 192.168.36.2:905
 route-target both 905:905
!
mac vrf eline906
 rd 192.168.36.2:906
 route-target both 906:906
!
mac vrf eline907
 rd 192.168.36.2:907
 route-target both 907:907
!
mac vrf eline908
 rd 192.168.36.2:908
 route-target both 908:908
!
mac vrf eline909
```

```
 rd 192.168.36.2:909
 route-target both 909:909
!
mac vrf eline910
 rd 192.168.36.2:910
 route-target both 910:910
!
mac vrf irbElan3001
 rd 192.168.36.2:1001
 route-target both 4294967294:3001
!
mac vrf irbElan3002
 rd 192.168.36.2:1002
 route-target export 4294967294:301
 route-target import 4294967294:302
!
mac vrf irbElan3003
 rd 192.168.36.2:1003
 route-target both 4294967294:3003
!
evpn mpls vtep-ip-global 192.168.36.2
!
evpn mpls mac-ageing-time 20
!
evpn mpls id 1901
 host-reachability-protocol evpn-bgp ELAN1901
!
evpn mpls id 61902
 host-reachability-protocol evpn-bgp ELAN1902
!
evpn mpls id 61904
 host-reachability-protocol evpn-bgp ELAN1904
!
evpn mpls id 163001
 host-reachability-protocol evpn-bgp irbElan3001
 evpn irb irb3001
!
evpn mpls id 163002
 host-reachability-protocol evpn-bgp irbElan3002
 evpn irb irb3002
!
evpn mpls id 163003
 host-reachability-protocol evpn-bgp irbElan3003
 evpn irb irb3003
!
evpn mpls id 1671905
 host-reachability-protocol evpn-bgp ELAN1905
!
evpn mpls id 1671906
 host-reachability-protocol evpn-bgp ELAN1906
!
evpn mpls id 1671907
 host-reachability-protocol evpn-bgp ELAN1907
!
evpn mpls id 1671908
 host-reachability-protocol evpn-bgp ELAN1908
!
evpn mpls id 1671909
 host-reachability-protocol evpn-bgp ELAN1909
!
evpn mpls id 1671910
 host-reachability-protocol evpn-bgp ELAN1910
!
evpn mpls id 16772901 xconnect target-mpls-id 16771901
 host-reachability-protocol evpn-bgp eline901
!
evpn mpls id 16772902 xconnect target-mpls-id 16771902
 host-reachability-protocol evpn-bgp eline902
```

© 2026 IP Infusion Inc. Proprietary

```
!
evpn mpls id 16772903 xconnect target-mpls-id 16771903
 host-reachability-protocol evpn-bgp eline903
!
evpn mpls id 16772904 xconnect target-mpls-id 16771904
 host-reachability-protocol evpn-bgp eline904
!
evpn mpls id 16772905 xconnect target-mpls-id 16771905
 host-reachability-protocol evpn-bgp eline905
!
evpn mpls id 16772906 xconnect target-mpls-id 16771906
 host-reachability-protocol evpn-bgp eline906
!
evpn mpls id 16772907 xconnect target-mpls-id 16771907
 host-reachability-protocol evpn-bgp eline907
!
evpn mpls id 16772908 xconnect target-mpls-id 16771908
 host-reachability-protocol evpn-bgp eline908
!
evpn mpls id 16772909 xconnect target-mpls-id 16771909
 host-reachability-protocol evpn-bgp eline909
!
evpn mpls id 16772910 xconnect target-mpls-id 16771910
 host-reachability-protocol evpn-bgp eline910
!
evpn mpls id 16777215
 host-reachability-protocol evpn-bgp ELAN1903
!
segment-routing
!
ip multicast-routing
!
ipv6 multicast-routing
!
ip prefix-list DEFAULT
 seq 10 permit 0.0.0.0/0
!
ip prefix-list LDP
 seq 5 deny any
!
ip prefix-list LOOPBACK
 seq 10 permit 192.168.36.2/32
!
ip prefix-list PL-CUST-SUBNETS-IRBVRF3003-IPv4
 seq 5 permit 192.2.2.0/24
!
ip prefix-list PL-DENY-DEFAULT-IRBVRF3003-IPv4
 seq 5 permit 0.0.0.0/0
!
ip prefix-list PL-DNS-SERVERS-IRBVRF3003-IPv4
 seq 5 permit 172.16.30.53/32
!
ip prefix-list PL-EXPORT-INTERVRFv4
 seq 5 deny 5.5.5.0/24
 seq 10 permit 5.5.6.0/24
 seq 11 permit 5.5.7.0/24
!
ip prefix-list PL-NEXTHOP-IRBVRF3003-IPv4
 seq 5 permit 80.12.1.254/32
 seq 10 permit 201.103.1.2/32
!
ip prefix-list PL-NOMETRIC-IRBVRF3003-IPv4
 seq 5 permit 192.168.30.0/24
!
mac-list PL-ELAN1901-MAC-HOSTS
 seq 10 permit 0010.9400.0002 0010.9400.0002
!
ipv6 prefix-list PL-CUST-SUBNETS-IRBVRF3003-IPv6
```

```
 seq 5 permit 2001:db8:3003:20::/64
!
ipv6 prefix-list PL-DENY-DEFAULT-IRBVRF3003-IPv6
 seq 5 permit ::/0
!
ipv6 prefix-list PL-DNS-SERVERS-IRBVRF3003-IPv6
 seq 5 permit 2001:db8:3003:53::53/128
!
ipv6 prefix-list PL-NEXTHOP-IRBVRF3003-IPv6
 seq 5 permit 80:12:1::254/128
 seq 10 permit 2001:bd8:103::1/128
!
ipv6 prefix-list PL-NOMETRIC-IRBVRF3003-IPv6
 seq 5 permit 2001:db8:3003:10::/64
!
ipv6 prefix-list PL-VRFLEAKINGv6
 seq 5 deny 2222:1:1:1::/64
 seq 6 permit 2222:1:1:2::/64
 seq 7 permit 2222:1:1:3::/64
!
router ldp
 router-id 192.168.36.2
 fast-reroute
 pw-status-tlv
 ignore-mac-withdraw-bad-pdu-length
 targeted-peer ipv4 192.168.36.1
  exit-targeted-peer-mode
 targeted-peer ipv4 192.168.36.3
  exit-targeted-peer-mode
 targeted-peer ipv4 192.168.36.4
  exit-targeted-peer-mode
 targeted-peer ipv4 192.168.36.11
  exit-targeted-peer-mode
 targeted-peer ipv4 192.168.36.12
  exit-targeted-peer-mode
 transport-address ipv4 192.168.36.2
 neighbor all tcp-mss 1440
 neighbor all auth md5 password plain-text P@ssw0rd
!
router rsvp
 lsp-reoptimization-timer 2
 hello-interval 3
 hello-timeout 11
 from 192.168.36.2
 detour-allow-primary-upstream-path
 detour-identification path
 entropy-label-capability
 auto-bypass
  attributes best-effort
   reoptimize
   exit
  inactivity-timer 30
  enable
  exit
 auto-bandwidth-on-boot 1 5 1
!
route-map REDISTRIBUTE-CONNECTED-TO-BGP permit 10
 match ip address prefix-list LOOPBACK
!
route-map RM-EXPORT-ELAN1901 permit 710
 match mac address list PL-ELAN1901-MAC-HOSTS
 set metric 300
 set local-preference 300
 set aigp-metric 100
 set atomic-aggregate
 set community 100:111
 set large-community 1:1:1
 set extcommunity rt 1:1 additive
```

```
!
route-map RM-EXPORT-ATTR permit 10
 match ip address prefix-list PL-EXPORT-INTERVRFv4
 set tag 1000
 set metric 4294967295
 set local-preference 4294967295
 set origin igp
 set aigp-metric 2000
 set as-path tag
 set community 100:1 additive
 set large-community 22:22:22
 set extcommunity rt 120:10 230:10 additive
!
route-map RM-EXPORT-ATTR permit 20
 match ipv6 address prefix-list PL-VRFLEAKINGv6
 set tag 1001
 set metric 4294967295
 set local-preference 4294967295
 set origin igp
 set aigp-metric 3000
 set as-path tag
 set community 100:1 additive
 set large-community 23:34:45
 set extcommunity rt 120:120 230:130 additive
!
route-map RM-EXPORT-ATTR permit 30
!
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1001
 match ip address prefix-list PL-NOMETRIC-IRBVRF3003-IPv4
 continue 1002
 set metric +1001
 set local-preference 1001
 set aigp-metric 1001
 set extcommunity rt 51185:1001 additive
!
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1002
 match ip address prefix-list PL-CUST-SUBNETS-IRBVRF3003-IPv4
 continue 1003
 set metric +1002
 set local-preference 1002
 set community 51185:1002 additive
 set large-community 4294967294:51185:1002 additive
!
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1003
 match ip address prefix-list PL-DENY-DEFAULT-IRBVRF3003-IPv4
 continue 1004
 set metric +1003
 set local-preference 1003
 set aigp-metric 1003
 set extcommunity rt 51185:1003 additive
!
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1004
 match ip address prefix-list PL-DNS-SERVERS-IRBVRF3003-IPv4
 continue 1005
 set tag 1001
 set extcommunity color 4200000001
 set atomic-aggregate
 set metric 2345
 set local-preference 4567
 set origin igp
 set aigp-metric 3000
 set as-path tag
 set community 4200000002
 set large-community 65000:400:40
 set aggregator as 65080 192.0.2.80
 set extcommunity rt 4200000001:20 additive
 set extcommunity cost 99 900
!
```

© 2026 IP Infusion Inc. Proprietary

```
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1005
 match ip next-hop prefix-list PL-NEXTHOP-IRBVRF3003-IPv4
 continue 1006
 set metric +1005
 set local-preference 1005
 set aigp-metric 1005
 set extcommunity rt 51185:1005 additive
!
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1006
 match community CUSTOMER_ROUTES_3015_IRBVRF3003_IPv4
 continue 1007
 set metric +1006
 set local-preference 1006
 set community 51185:1006 additive
 set large-community 4294967294:51185:1006 additive
!
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1007
 match extcommunity CUSTOMER_EXTENDED_COMM_3016_IRBVRF3003_IPv4
 continue 1008
 set metric +1007
 set local-preference 1007
 set aigp-metric 1007
 set extcommunity rt 51185:1007 additive
!
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1008
 match extcommunity CUSTOMER_EXTENDED_COMM_3017_4byte_IRBVRF3003_IPv4
 continue 1009
 set metric +1008
 set local-preference 1008
 set community 51185:1008 additive
 set large-community 4294967294:51185:1008 additive
!
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1009
 match large-community CUSTOMER_LARGE_COMM_IRBVRF3003_V4
 continue 1010
 set metric +1009
 set local-preference 1009
 set aigp-metric 1009
 set extcommunity rt 51185:1009 additive
!
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1010
 match as-path ASPATH-IRBVRF3003-IN-V4
 continue 1011
 set metric +1010
 set local-preference 1010
 set community 51185:1010 additive
 set large-community 4294967294:51185:1010 additive
!
interface cd48
!
interface cd49
!
interface cd50
!
interface cd51
!
interface cd52
!
interface cd53
!
interface cd54
!
interface cd55
!
interface ce1
!
interface ce2
!
```

```
interface ce3
!
interface ce4
!
interface ce5
!
interface ce6
!
interface ce7
!
interface ce8
!
interface ce9
!
interface ce10
!
interface ce11
!
interface ce12
!
interface ce13
!
interface ce14
!
interface ce15
!
interface ce16
!
interface ce17
!
interface ce18
!
interface ce19
!
interface ce20
!
interface ce21
!
interface ce22
!
interface ce23
!
interface ce24
!
interface ce25
!
interface ce26
!
interface ce27
!
interface ce28
!
interface ce29
!
interface ce30
!
interface ce31
!
interface ce32
!
interface ce33
!
interface ce34
!
interface ce35
!
interface ce36
!
```

© 2026 IP Infusion Inc. Proprietary

```
interface ce37
!
interface ce38
!
interface ce39
!
interface ce40
!
interface ce41
!
interface ce42
!
interface ce43
!
interface ce44
!
interface ce45
!
interface ce46/1
!
interface ce46/2
 description connected_to_PE3
 load-interval 30
 ip address 203.0.113.18/31
 ipv6 address 203:3:8::105/64
 mtu 9194
 label-switching
 link-debounce-time 2000 0
 mpls ldp-igp sync isis level-2 holddown-timer 900
 isis network point-to-point
 ip router isis 1
 ipv6 router isis 1
 isis authentication mode md5 level-1
 isis authentication mode md5 level-2
 isis authentication key-chain isis level-1
 isis authentication key-chain isis level-2
 mpls ldp-igp sync-delay 30
 enable-rsvp
 ip pim sparse-mode
 lldp-agent
  set lldp enable txrx
  set lldp chassis-id-tlv ip-address
  set lldp port-id-tlv if-name
  lldp tlv basic-mgmt system-name select
  lldp tlv basic-mgmt system-description select
  exit
 bfd interval 3 minrx 3 multiplier 3
!
interface ce46/3
 description ### Link to RR2 ##
 load-interval 30
 mtu 9198
 link-debounce-time 2000 0
 lldp-agent
  set lldp enable txrx
  set lldp chassis-id-tlv ip-address
  set lldp port-id-tlv if-name
  lldp tlv basic-mgmt system-name select
  lldp tlv basic-mgmt system-description select
  exit
!
interface ce46/3.101
 description ### Link to rr-2 ##
 encapsulation dot1q 101
 load-interval 30
 ip address 203.0.113.11/31
 ipv6 address 203:3:6::105/64
 mtu 9194
```

```
 label-switching
 mpls ldp-igp sync isis level-2 holddown-timer 900
 isis network point-to-point
 ip router isis 1
 ipv6 router isis 1
 isis authentication mode md5 level-1
 isis authentication mode md5 level-2
 isis authentication key-chain isis level-1
 isis authentication key-chain isis level-2
 enable-ldp ipv4
 mpls ldp-igp sync-delay 30
 enable-rsvp
 ip pim sparse-mode
 bfd interval 3 minrx 3 multiplier 3
!
interface ce46/4
 description connected_to_rr1
 load-interval 30
 ip address 203.0.113.16/31
 ipv6 address 203:3:7::105/64
 mtu 9194
 label-switching
 link-debounce-time 2000 0
 mpls ldp-igp sync isis level-2 holddown-timer 900
 isis network point-to-point
 ip router isis 1
 ipv6 router isis 1
 isis authentication mode md5 level-1
 isis authentication mode md5 level-2
 isis authentication key-chain isis level-1
 isis authentication key-chain isis level-2
 mpls ldp-igp sync-delay 30
 enable-rsvp
 ip pim sparse-mode
 lldp-agent
  set lldp enable txrx
  set lldp chassis-id-tlv ip-address
  set lldp port-id-tlv if-name
  lldp tlv basic-mgmt system-name select
  lldp tlv basic-mgmt system-description select
  exit
 bfd interval 3 minrx 3 multiplier 3
!
interface ce47/1
!
interface ce47/2
!
interface ce47/3
!
interface ce47/4
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface irb3001
 ip vrf forwarding irbvrf3001
 ip address 80.10.1.1/24
 ipv6 address 80:10:1::1/64
 mtu 9216
!
interface irb3002
 ip vrf forwarding irbvrf3002
 ip address 80.11.1.1/24
 ipv6 address 80:11:1::1/64
 mtu 9216
!
interface irb3003
```

```
 ip vrf forwarding irbvrf3003
 ip address 80.12.1.1/24
 ipv6 address 80:12:1::1/64
 mtu 9216
!
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 loopback1
 ip address 192.168.36.2/32
 ipv6 address cafe:168:36::2/128
 prefix-sid index 2 explicit-null
 ip router isis 1
 ipv6 router isis 1
 ip pim sparse-mode
!
interface xe0
!
interface xe1
 description ## Connected to Spirent-2/16 ##
 load-interval 30
 mtu 9216
!
interface xe1.10 switchport
 encapsulation dot1q 10
 load-interval 30
 mtu 8000
 access-if-vpws
  mpls-l2-circuit vc2001 primary
  mpls-l2-circuit vc2000 secondary
!
interface xe1.11 switchport
 encapsulation dot1q 11
 load-interval 30
 mtu 8000
 access-if-vpws
  mpls-l2-circuit vc2003 primary
  mpls-l2-circuit vc2004 secondary
!
interface xe1.101
 encapsulation dot1q 101
 load-interval 30
 ip vrf forwarding vrf101
 ip address 201.101.1.1/24
 ipv6 address 2001:bd8:101::1/64
 mtu 9216
!
interface xe1.102
 encapsulation dot1q 102
 load-interval 30
 ip vrf forwarding vrf102
 ip address 201.102.1.1/24
 ipv6 address 2001:bd8:102::1/64
 mtu 9216
!
interface xe1.103
 encapsulation dot1q 103
 load-interval 30
 ip vrf forwarding vrf103
 ip address 201.103.1.1/24
 ipv6 address 2001:bd8:103::1/64
 mtu 9216
```

© 2026 IP Infusion Inc. Proprietary

```
!
interface xe1.104
 encapsulation dot1q 104
 load-interval 30
 ip vrf forwarding vrf104
 ip address 201.104.1.1/24
 ipv6 address 2001:bd8:104::1/64
 mtu 9216
!
interface xe1.105
 encapsulation dot1q 105
 load-interval 30
 ip vrf forwarding vrf105
 ip address 201.105.1.1/24
 ipv6 address 2001:bd8:105::1/64
 mtu 9216
!
interface xe1.106
 encapsulation dot1q 106
 load-interval 30
 ip vrf forwarding vrf106
 ip address 201.106.1.1/24
 ipv6 address 2001:bd8:106::1/64
 mtu 9216
!
interface xe1.107
 encapsulation dot1q 107
 load-interval 30
 ip vrf forwarding vrf107
 ip address 201.107.1.1/24
 ipv6 address 2001:bd8:107::1/64
 mtu 9216
!
interface xe1.108
 encapsulation dot1q 108
 load-interval 30
 ip vrf forwarding vrf108
 ip address 201.108.1.1/24
 ipv6 address 2001:bd8:108::1/64
 mtu 9216
!
interface xe1.109
 encapsulation dot1q 109
 load-interval 30
 ip vrf forwarding vrf109
 ip address 201.109.1.1/24
 ipv6 address 2001:bd8:109::1/64
 mtu 9216
!
interface xe1.110
 encapsulation dot1q 110
 load-interval 30
 ip vrf forwarding vrf110
 ip address 201.110.1.1/24
 ipv6 address 2001:bd8:110::1/64
 mtu 9216
!
interface xe1.111
 encapsulation dot1q 111
 load-interval 30
 ip vrf forwarding vrf111
 ip address 201.111.1.1/24
 ipv6 address 2001:bd8:111::1/64
 mtu 9216
!
interface xe1.112
 encapsulation dot1q 112
 load-interval 30
```

```
 ip vrf forwarding vrf111
 ip address 201.112.1.1/24
 ipv6 address 2001:bd8:112::1/64
 mtu 9216
!
interface xe1.113
 encapsulation dot1q 113
 load-interval 30
 ip vrf forwarding vrf111
 ip address 201.113.1.1/24
 ipv6 address 2001:bd8:113::1/64
 mtu 9216
!
interface xe1.114
 encapsulation dot1q 114
 load-interval 30
 ip vrf forwarding vrf111
 ip address 201.114.1.1/24
 ipv6 address 2001:bd8:114::1/64
 mtu 9216
!
interface xe1.115
 encapsulation dot1q 115
 load-interval 30
 ip vrf forwarding vrf111
 ip address 201.115.1.1/24
 ipv6 address 2001:bd8:115::1/64
 mtu 9216
!
interface xe1.116
 encapsulation dot1q 116
 load-interval 30
 ip vrf forwarding vrf111
 ip address 201.116.1.1/24
 ipv6 address 2001:bd8:116::1/64
 mtu 9216
!
interface xe1.117
 encapsulation dot1q 117
 load-interval 30
 ip vrf forwarding vrf111
 ip address 201.117.1.1/24
 ipv6 address 2001:bd8:117::1/64
 mtu 9216
!
interface xe1.118
 encapsulation dot1q 118
 load-interval 30
 ip vrf forwarding vrf111
 ip address 201.118.1.1/24
 ipv6 address 2001:bd8:118::1/64
 mtu 9216
!
interface xe1.119
 encapsulation dot1q 119
 load-interval 30
 ip vrf forwarding vrf111
 ip address 201.119.1.1/24
 ipv6 address 2001:bd8:119::1/64
 mtu 9216
!
interface xe1.120
 encapsulation dot1q 120
 load-interval 30
 ip vrf forwarding vrf111
 ip address 201.120.1.1/24
 ipv6 address 2001:bd8:11a::1/64
 mtu 9216
```

© 2026 IP Infusion Inc. Proprietary

```
!
interface xe1.300
!
interface xe1.890
 description for ### IPv4 eBGP
 encapsulation dot1q 890
 load-interval 30
 ip address 190.160.2.1/24
 mtu 9216
!
interface xe1.891
 description for ### 6PE
 encapsulation dot1q 891
 load-interval 30
 ipv6 address 3601::1/64
 mtu 9216
!
interface xe1.892
 description for ### 6PE
 encapsulation dot1q 892
 load-interval 30
 ipv6 address 3602::1/64
 mtu 9216
!
interface xe1.893
 description for ### 6PE
 encapsulation dot1q 893
 load-interval 30
 ipv6 address 3603::1/64
 mtu 9216
!
interface xe1.894
 description for ### 6PE
 encapsulation dot1q 894
 load-interval 30
 ipv6 address 3604::1/64
 mtu 9216
!
interface xe1.895
 description for ### 6PE
 encapsulation dot1q 895
 load-interval 30
 ipv6 address 3605::1/64
 mtu 9216
!
interface xe1.896
 description for ### 6PE
 encapsulation dot1q 896
 load-interval 30
 ipv6 address 3606::1/64
 mtu 9216
!
interface xe1.897
 description for ### 6PE
 encapsulation dot1q 897
 load-interval 30
 ipv6 address 3607::1/64
 mtu 9216
!
interface xe1.898
 description for ### 6PE
 encapsulation dot1q 898
 load-interval 30
 ipv6 address 3608::1/64
 mtu 9216
!
interface xe1.899
 description for ### 6PE
```

```
 encapsulation dot1q 899
 load-interval 30
 ipv6 address 3609::1/64
 mtu 9216
!
interface xe1.900
 description for ### 6PE
 encapsulation dot1q 900
 load-interval 30
 ipv6 address 360a::1/64
 mtu 9216
!
interface xe1.901 switchport
 description for ### eline901
 encapsulation dot1q 901
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 16772901
!
interface xe1.902 switchport
 description for ### eline902
 encapsulation dot1q 902
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 16772902
!
interface xe1.903 switchport
 description for ### eline903
 encapsulation dot1q 903
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 16772903
!
interface xe1.904 switchport
 description for ### eline904
 encapsulation dot1q 904
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 16772904
!
interface xe1.905 switchport
 description for ### eline905
 encapsulation dot1q 905
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 16772905
!
interface xe1.906 switchport
 description for ### eline906
 encapsulation dot1q 906
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 16772906
!
interface xe1.907 switchport
 description for ### eline907
 encapsulation dot1q 907
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 16772907
!
```

```
interface xe1.908 switchport
 description for ### eline908
 encapsulation dot1q 908
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 16772908
!
interface xe1.909 switchport
 description for ### eline909
 encapsulation dot1q 909
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 16772909
!
interface xe1.910 switchport
 description for ### eline910
 encapsulation dot1q 910
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 16772910
!
interface xe1.1250 switchport
 description for ### bgp vpls1250
 encapsulation dot1q 1250
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls4294961250
!
interface xe1.1251 switchport
 description for ### bgp vpls1251
 encapsulation dot1q 1251
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls4294961251
!
interface xe1.1252 switchport
 description for ### bgp vpls1252
 encapsulation dot1q 1252
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls4294961252
!
interface xe1.1253 switchport
 description for ### bgp vpls1253
 encapsulation dot1q 1253
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls4294961253
!
interface xe1.1254 switchport
 description for ### bgp vpls1254
 encapsulation dot1q 1254
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls4294961254
!
interface xe1.1255 switchport
 description for ### bgp vpls1255
 encapsulation dot1q 1255
 load-interval 30
```

```
 mtu 9216
 access-if-vpls
  mpls-vpls vpls4294961255
!
interface xe1.1256 switchport
 description for ### bgp vpls1256
 encapsulation dot1q 1256
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls4294961256
!
interface xe1.1257 switchport
 description for ### bgp vpls1257
 encapsulation dot1q 1257
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls4294961257
!
interface xe1.1258 switchport
 description for ### bgp vpls1258
 encapsulation dot1q 1258
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls4294961258
!
interface xe1.1259 switchport
 description for ### bgp vpls1259
 encapsulation dot1q 1259
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls4294961259
!
interface xe1.1260 switchport
 description for ### bgp vpls1260
 encapsulation dot1q 1260
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls4294961260
!
interface xe1.1261 switchport
 description for ### LDP_vpls1261
 encapsulation dot1q 1261
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls1261
!
interface xe1.1262 switchport
 description for ### LDP_vpls1262
 encapsulation dot1q 1262
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls1262
!
interface xe1.1263 switchport
 description for ### LDP_vpls1263
 encapsulation dot1q 1263
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls1263
!
```

```
interface xe1.1264 switchport
 description for ### LDP_vpls1264
 encapsulation dot1q 1264
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls1264
!
interface xe1.1265 switchport
 description for ### LDP_vpls1265
 encapsulation dot1q 1265
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls1265
!
interface xe1.1266 switchport
 description for ### LDP_vpls1266
 encapsulation dot1q 1266
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls1266
!
interface xe1.1267 switchport
 description for ### LDP_vpls1267
 encapsulation dot1q 1267
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls1267
!
interface xe1.1268 switchport
 description for ### LDP_vpls1268
 encapsulation dot1q 1268
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls1268
!
interface xe1.1269 switchport
 description for ### LDP_vpls1269
 encapsulation dot1q 1269
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls1269
!
interface xe1.1270 switchport
 description for ### LDP_vpls1270
 encapsulation dot1q 1270
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls vpls1270
!
interface xe1.1271 switchport
 description for ### mpls-l2-circuit-1271
 encapsulation dot1q 1271
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1271 primary
!
interface xe1.1272 switchport
 description for ### mpls-l2-circuit-1272
 encapsulation dot1q 1272
 load-interval 30
```

```
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1272 primary
!
interface xe1.1273 switchport
 description for ### mpls-l2-circuit-1273
 encapsulation dot1q 1273
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1273 primary
!
interface xe1.1274 switchport
 description for ### mpls-l2-circuit-1274
 encapsulation dot1q 1274
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1274 primary
!
interface xe1.1275 switchport
 description for ### mpls-l2-circuit-1275
 encapsulation dot1q 1275
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1275 primary
!
interface xe1.1276 switchport
 description for ### mpls-l2-circuit-1276
 encapsulation dot1q 1276
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1276 primary
!
interface xe1.1277 switchport
 description for ### mpls-l2-circuit-1277
 encapsulation dot1q 1277
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1277 primary
!
interface xe1.1278 switchport
 description for ### mpls-l2-circuit-1278
 encapsulation dot1q 1278
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1278 primary
!
interface xe1.1279 switchport
 description for ### mpls-l2-circuit-1279
 encapsulation dot1q 1279
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1279 primary
!
interface xe1.1280 switchport
 description for ### mpls-l2-circuit-1280
 encapsulation dot1q 1280
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1280 primary
!
```

© 2026 IP Infusion Inc. Proprietary

```
interface xe1.1281 switchport
 description for ### mpls-l2-circuit-1281
 encapsulation dot1q 1281
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1281 primary
!
interface xe1.1282 switchport
 description for ### mpls-l2-circuit-1282
 encapsulation dot1q 1282
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1282 primary
!
interface xe1.1283 switchport
 description for ### mpls-l2-circuit-1283
 encapsulation dot1q 1283
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1283 primary
!
interface xe1.1284 switchport
 description for ### mpls-l2-circuit-1284
 encapsulation dot1q 1284
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1284 primary
!
interface xe1.1285 switchport
 description for ### mpls-l2-circuit-1285
 encapsulation dot1q 1285
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1285 primary
!
interface xe1.1286 switchport
 description for ### mpls-l2-circuit-1286
 encapsulation dot1q 1286
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1286 primary
!
interface xe1.1287 switchport
 description for ### mpls-l2-circuit-1287
 encapsulation dot1q 1287
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1287 primary
!
interface xe1.1288 switchport
 description for ### mpls-l2-circuit-1288
 encapsulation dot1q 1288
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1288 primary
!
interface xe1.1289 switchport
 description for ### mpls-l2-circuit-1289
 encapsulation dot1q 1289
 load-interval 30
```

```
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1289 primary
!
interface xe1.1290 switchport
 description for ### mpls-l2-circuit-1290
 encapsulation dot1q 1290
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1290 primary
!
interface xe1.1291 switchport
 description for ### mpls-l2-circuit-1291
 encapsulation dot1q 1291
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1291 primary
!
interface xe1.1292 switchport
 description for ### mpls-l2-circuit-1292
 encapsulation dot1q 1292
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1292 primary
!
interface xe1.1293 switchport
 description for ### mpls-l2-circuit-1293
 encapsulation dot1q 1293
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1293 primary
!
interface xe1.1294 switchport
 description for ### mpls-l2-circuit-1294
 encapsulation dot1q 1294
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1294 primary
!
interface xe1.1295 switchport
 description for ### mpls-l2-circuit-1295
 encapsulation dot1q 1295
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1295 primary
!
interface xe1.1296 switchport
 description for ### mpls-l2-circuit-1296
 encapsulation dot1q 1296
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1296 primary
!
interface xe1.1297 switchport
 description for ### mpls-l2-circuit-1297
 encapsulation dot1q 1297
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1297 primary
!
```

```
interface xe1.1298 switchport
 description for ### mpls-l2-circuit-1298
 encapsulation dot1q 1298
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1298 primary
!
interface xe1.1299 switchport
 description for ### mpls-l2-circuit-1299
 encapsulation dot1q 1299
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1299 primary
!
interface xe1.1300 switchport
 description for ### mpls-l2-circuit-1300
 encapsulation dot1q 1300
 load-interval 30
 mtu 9216
 access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1300 primary
!
interface xe1.1301 switchport
 description for ### bgp-ad-1301
 encapsulation dot1q 1301
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls VPLS_AD1301
!
interface xe1.1302 switchport
 description for ### bgp-ad-1302
 encapsulation dot1q 1302
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls VPLS_AD1302
!
interface xe1.1303 switchport
 description for ### bgp-ad-1303
 encapsulation dot1q 1303
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls VPLS_AD1303
!
interface xe1.1304 switchport
 description for ### bgp-ad-1304
 encapsulation dot1q 1304
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls VPLS_AD1304
!
interface xe1.1305 switchport
 description for ### bgp-ad-1305
 encapsulation dot1q 1305
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls VPLS_AD1305
!
interface xe1.1306 switchport
 description for ### bgp-ad-1306
 encapsulation dot1q 1306
 load-interval 30
```

© 2026 IP Infusion Inc. Proprietary

```
 mtu 9216
 access-if-vpls
  mpls-vpls VPLS_AD1306
!
interface xe1.1307 switchport
 description for ### bgp-ad-1307
 encapsulation dot1q 1307
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls VPLS_AD1307
!
interface xe1.1308 switchport
 description for ### bgp-ad-1308
 encapsulation dot1q 1308
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls VPLS_AD1308
!
interface xe1.1309 switchport
 description for ### bgp-ad-1309
 encapsulation dot1q 1309
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls VPLS_AD1309
!
interface xe1.1310 switchport
 description for ### bgp-ad-1310
 encapsulation dot1q 1310
 load-interval 30
 mtu 9216
 access-if-vpls
  mpls-vpls VPLS_AD1310
!
interface xe1.1901 switchport
 description for ### elan1901
 encapsulation dot1q 1901
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 1901
!
interface xe1.1902 switchport
 description for ### elan1902
 encapsulation dot1q 1902
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 61902
!
interface xe1.1903 switchport
 description for ### elan1903
 encapsulation dot1q 1903
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 16777215
!
interface xe1.1904 switchport
 description for ### elan1904
 encapsulation dot1q 1904
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 61904
!
```

```
interface xe1.1905 switchport
 description for ### elan1905
 encapsulation dot1q 1905
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 1671905
!
interface xe1.1906 switchport
 description for ### elan1906
 encapsulation dot1q 1906
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 1671906
!
interface xe1.1907 switchport
 description for ### elan1907
 encapsulation dot1q 1907
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 1671907
!
interface xe1.1908 switchport
 description for ### elan1908
 encapsulation dot1q 1908
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 1671908
!
interface xe1.1909 switchport
 description for ### elan1909
 encapsulation dot1q 1909
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 1671909
!
interface xe1.1910 switchport
 description for ### elan1910
 encapsulation dot1q 1910
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 1671910
!
interface xe1.3001 switchport
 description for ### irbvrf3001
 encapsulation dot1q 3001
 rewrite pop
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 163001
!
interface xe1.3002 switchport
 description for ### irbvrf3002
 encapsulation dot1q 3002
 rewrite pop
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 163002
!
interface xe1.3003 switchport
 description for ### irbvrf3003
```

```
 encapsulation dot1q 3003
 rewrite pop
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 163003
!
interface xe2
!
interface xe3
!
 exit
!
router ospf 65535
 fast-reroute keep-all-paths
 shutdown
 bfd all-interfaces
 fast-reroute per-prefix remote-lfa area 0.0.0.1 tunnel mpls-ldp
 network 192.168.36.2/32 area 0.0.0.1
 network 203.0.113.10/31 area 0.0.0.1
 network 203.0.113.16/31 area 0.0.0.1
 network 203.0.113.18/31 area 0.0.0.1
!
router isis 1
 is-type level-2-only
 authentication mode md5 level-2
 authentication key-chain isis level-2
 ignore-lsp-errors
 lsp-gen-interval 5
 max-lsp-lifetime 2000
 spf-interval-exp level-2 50 2000
 metric-style wide
 microloop-avoidance level-2
 microloop-avoidance max-fib 60 level-2
 mpls traffic-eng router-id 192.168.36.2
 mpls traffic-eng level-2
 capability cspf
 dynamic-hostname
 fast-reroute terminate-hold-on interval 100000
 fast-reroute per-prefix level-2 proto ipv4 all
 fast-reroute per-prefix remote-lfa level-2 proto ipv4 tunnel mpls-ldp
 fast-reroute ti-lfa level-2 proto ipv4
 bfd all-interfaces
 net 49.0000.0000.0002.00
 passive-interface loopback1
 isis segment-routing global block 16000 23999
 segment-routing entropy-label
!
router bgp 4294967294
 bgp router-id 192.168.36.2
 bgp auto-policy-soft-reset enable
 bgp log-neighbor-changes
 no bgp inbound-route-filter
 allocate-label all
 neighbor PEER-BGPLU peer-group
 neighbor PEER-BGPLU remote-as 2200000002
 neighbor PEER-RR peer-group
 neighbor PEER-RR remote-as 4294967294
 neighbor PEER-RR tcp-mss 1440
 neighbor PEER-RR update-source loopback1
 neighbor PEER-RR authentication-key
0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157be7f6e35f32bf637
 neighbor PEER-RR advertisement-interval 0
 neighbor PEER-RR fall-over bfd multihop
 neighbor 190.160.2.254 peer-group PEER-BGPLU
 neighbor 192.168.36.11 remote-as 4294967294
 neighbor 192.168.36.11 tcp-mss 1440
 neighbor 192.168.36.11 update-source loopback1
 neighbor 192.168.36.11 authentication-key
```

```
0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157be7f6e35f32bf637
 neighbor 192.168.36.11 advertisement-interval 0
 neighbor 192.168.36.11 fall-over bfd multihop
 neighbor 192.168.36.12 peer-group PEER-RR
 neighbor 3601::2 remote-as 200
 neighbor 3602::2 remote-as 200
 neighbor 3603::2 remote-as 200
 neighbor 3604::2 remote-as 200
 neighbor 3605::2 remote-as 200
 neighbor 3606::2 remote-as 200
 neighbor 3607::2 remote-as 200
 neighbor 3608::2 remote-as 200
 neighbor 3609::2 remote-as 200
 neighbor 360a::2 remote-as 200
 !
 address-family ipv4 unicast
 redistribute connected route-map REDISTRIBUTE-CONNECTED-TO-BGP
 neighbor PEER-RR activate
 neighbor 192.168.36.11 activate
 exit-address-family
 !
 address-family ipv4 labeled-unicast
 neighbor PEER-BGPLU activate
 neighbor PEER-RR activate
 neighbor 192.168.36.11 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor PEER-RR activate
 neighbor 192.168.36.11 activate
 exit-address-family
 !
 address-family rtfilter unicast
 exit-address-family
 !
 address-family l2vpn vpls
 neighbor PEER-RR activate
 neighbor 192.168.36.11 activate
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor PEER-RR activate
 neighbor 192.168.36.11 activate
 exit-address-family
 !
 address-family vpnv6 unicast
 neighbor PEER-RR activate
 neighbor 192.168.36.11 activate
 exit-address-family
 !
 address-family ipv6 unicast
 redistribute connected
 neighbor 3601::2 activate
 neighbor 3602::2 activate
 neighbor 3603::2 activate
 neighbor 3604::2 activate
 neighbor 3605::2 activate
 neighbor 3606::2 activate
 neighbor 3607::2 activate
 neighbor 3608::2 activate
 neighbor 3609::2 activate
 neighbor 360a::2 activate
 exit-address-family
 !
 address-family ipv6 labeled-unicast
 neighbor PEER-RR activate
 neighbor 192.168.36.11 activate
 exit-address-family
```

```
!
address-family ipv4 vrf irbvrf3001
redistribute connected
neighbor 80.10.1.254 remote-as 65535
neighbor 80.10.1.254 activate
neighbor 80.10.1.254 authentication-key 0x503653bfffeb7c928057183d8be815ab
exit-address-family
!
address-family ipv4 vrf irbvrf3002
redistribute connected
neighbor 80.11.1.254 remote-as 3002
neighbor 80.11.1.254 activate
neighbor 80.11.1.254 authentication-key 0x503653bfffeb7c928057183d8be815ab
exit-address-family
!
address-family ipv4 vrf irbvrf3003
redistribute connected
neighbor 80.12.1.254 remote-as 3003
neighbor 80.12.1.254 activate
neighbor 80.12.1.254 authentication-key 0x503653bfffeb7c928057183d8be815ab
exit-address-family
!
address-family ipv4 vrf vrf101
redistribute connected
redistribute static
neighbor CLIENTS-V4 peer-group
neighbor CLIENTS-V4 remote-as 65535
neighbor CLIENTS-V4 activate
neighbor CLIENTS-V4 authentication-key 0x503653bfffeb7c928057183d8be815ab
neighbor CLIENTS-V4 ebgp-multihop 255
neighbor 201.101.1.2 peer-group CLIENTS-V4
exit-address-family
!
address-family ipv4 vrf vrf102
redistribute connected
neighbor 201.102.1.2 remote-as 65535
neighbor 201.102.1.2 activate
exit-address-family
!
address-family ipv4 vrf vrf103
redistribute connected
neighbor 201.103.1.2 remote-as 65535
neighbor 201.103.1.2 activate
exit-address-family
!
address-family ipv4 vrf vrf104
redistribute connected
neighbor 201.104.1.2 remote-as 65535
neighbor 201.104.1.2 activate
exit-address-family
!
address-family ipv4 vrf vrf105
redistribute connected
neighbor 201.105.1.2 remote-as 65535
neighbor 201.105.1.2 activate
exit-address-family
!
address-family ipv4 vrf vrf106
redistribute connected
neighbor 201.106.1.2 remote-as 65535
neighbor 201.106.1.2 activate
exit-address-family
!
address-family ipv4 vrf vrf107
redistribute connected
neighbor 201.107.1.2 remote-as 65535
neighbor 201.107.1.2 activate
exit-address-family
```

```
!
address-family ipv4 vrf vrf108
redistribute connected
neighbor 201.108.1.2 remote-as 65535
neighbor 201.108.1.2 activate
exit-address-family
!
address-family ipv4 vrf vrf109
redistribute connected
neighbor 201.109.1.2 remote-as 65535
neighbor 201.109.1.2 activate
exit-address-family
!
address-family ipv4 vrf vrf110
redistribute connected
neighbor 201.110.1.2 remote-as 65535
neighbor 201.110.1.2 activate
exit-address-family
!
address-family ipv4 vrf vrf111
neighbor CLIENTS peer-group
neighbor CLIENTS remote-as 65534
neighbor CLIENTS activate
neighbor CLIENTS ebgp-multihop 2
neighbor 201.111.1.2 remote-as 65535
neighbor 201.111.1.2 activate
neighbor 201.112.1.2 peer-group CLIENTS
neighbor 201.113.1.2 remote-as 65533
neighbor 201.113.1.2 activate
neighbor 201.114.1.2 remote-as 65532
neighbor 201.114.1.2 activate
neighbor 201.115.1.2 remote-as 65531
neighbor 201.115.1.2 activate
neighbor 201.116.1.2 remote-as 65530
neighbor 201.116.1.2 activate
neighbor 201.117.1.2 remote-as 65529
neighbor 201.117.1.2 activate
neighbor 201.118.1.2 remote-as 65528
neighbor 201.118.1.2 activate
neighbor 201.119.1.2 remote-as 65527
neighbor 201.119.1.2 activate
neighbor 201.120.1.2 remote-as 65526
neighbor 201.120.1.2 activate
exit-address-family
!
address-family ipv6 vrf irbvrf3001
redistribute connected
neighbor 80:10:1::254 remote-as 65535
neighbor 80:10:1::254 activate
neighbor 80:10:1::254 authentication-key 0x503653bfffeb7c928057183d8be815ab
exit-address-family
!
address-family ipv6 vrf irbvrf3002
redistribute connected
neighbor 80:11:1::254 remote-as 3002
neighbor 80:11:1::254 activate
neighbor 80:11:1::254 authentication-key 0x503653bfffeb7c928057183d8be815ab
exit-address-family
!
address-family ipv6 vrf irbvrf3003
redistribute connected
neighbor 80:12:1::254 remote-as 3003
neighbor 80:12:1::254 activate
neighbor 80:12:1::254 authentication-key 0x503653bfffeb7c928057183d8be815ab
exit-address-family
!
address-family ipv6 vrf vrf101
redistribute connected
```

© 2026 IP Infusion Inc. Proprietary

```
neighbor CLIENTS peer-group
neighbor CLIENTS remote-as 65535
neighbor CLIENTS activate
neighbor CLIENTS authentication-key 0x503653bfffeb7c928057183d8be815ab
neighbor CLIENTS ebgp-multihop 255
neighbor 2001:bd8:101::2 peer-group CLIENTS
exit-address-family
!
address-family ipv6 vrf vrf102
redistribute connected
neighbor 2001:bd8:102::2 remote-as 65535
neighbor 2001:bd8:102::2 activate
exit-address-family
!
address-family ipv6 vrf vrf103
redistribute connected
neighbor 2001:bd8:103::2 remote-as 65535
neighbor 2001:bd8:103::2 activate
exit-address-family
!
address-family ipv6 vrf vrf104
redistribute connected
neighbor 2001:bd8:104::2 remote-as 65535
neighbor 2001:bd8:104::2 activate
exit-address-family
!
address-family ipv6 vrf vrf105
redistribute connected
neighbor 2001:bd8:105::2 remote-as 65535
neighbor 2001:bd8:105::2 activate
exit-address-family
!
address-family ipv6 vrf vrf106
redistribute connected
neighbor 2001:bd8:106::2 remote-as 65535
neighbor 2001:bd8:106::2 activate
exit-address-family
!
address-family ipv6 vrf vrf107
redistribute connected
neighbor 2001:bd8:107::2 remote-as 65535
neighbor 2001:bd8:107::2 activate
exit-address-family
!
address-family ipv6 vrf vrf108
redistribute connected
neighbor 2001:bd8:108::2 remote-as 65535
neighbor 2001:bd8:108::2 activate
exit-address-family
!
address-family ipv6 vrf vrf109
redistribute connected
neighbor 2001:bd8:109::2 remote-as 65535
neighbor 2001:bd8:109::2 activate
exit-address-family
!
address-family ipv6 vrf vrf110
redistribute connected
neighbor 2001:bd8:110::2 remote-as 65535
neighbor 2001:bd8:110::2 activate
exit-address-family
!
address-family ipv6 vrf vrf111
redistribute connected
neighbor 2001:bd8:111::254 remote-as 65535
neighbor 2001:bd8:111::254 activate
neighbor 2001:bd8:112::254 remote-as 65534
neighbor 2001:bd8:112::254 activate
```

```
 neighbor 2001:bd8:113::254 remote-as 65533
 neighbor 2001:bd8:113::254 activate
 neighbor 2001:bd8:114::254 remote-as 65532
 neighbor 2001:bd8:114::254 activate
 neighbor 2001:bd8:115::254 remote-as 65531
 neighbor 2001:bd8:115::254 activate
 neighbor 2001:bd8:116::254 remote-as 65530
 neighbor 2001:bd8:116::254 activate
 neighbor 2001:bd8:117::254 remote-as 65529
 neighbor 2001:bd8:117::254 activate
 neighbor 2001:bd8:118::254 remote-as 65528
 neighbor 2001:bd8:118::254 activate
 neighbor 2001:bd8:119::254 remote-as 65527
 neighbor 2001:bd8:119::254 activate
 neighbor 2001:bd8:11a::254 remote-as 65526
 neighbor 2001:bd8:11a::254 activate
 exit-address-family
 !
 exit
!
rsvp-trunk PE2_1_to_RR_1 ipv4
 reoptimize
 primary fast-reroute protection facility
 primary fast-reroute node-protection
 update-type make-before-break
 to 192.168.36.12
!
rsvp-trunk PE2_1_to_RR2_1 ipv4
 reoptimize
 primary fast-reroute protection one-to-one
 primary fast-reroute node-protection
 primary label-record
 update-type make-before-break
 to 192.168.36.11
!
ip route vrf vrf101 0.0.0.0/0 Null
!
ip community-list standard CUSTOMER_ROUTES_3015_IRBVRF3003_IPv4 permit 51185:1015
ip community-list standard CUSTOMER_ROUTES_3015_IRBVRF3003_IPv6 permit 51185:3015
ip community-list standard CUSTOMER_ROUTES_ELAN1901 permit 1901:1901
!
ip large-community-list standard CUSTOMER_LARGE_COMM_ELAN1901 permit 1901:1901:65000
ip large-community-list standard CUSTOMER_LARGE_COMM_IRBVRF3003_V4 permit 200:3003:65000
ip large-community-list standard CUSTOMER_LARGE_COMM_IRBVRF3003_V6 permit 200:112:65000
!
ip extcommunity-list standard CUSTOMER_EXTENDED_COMM_3016_IRBVRF3003_IPv4 permit rt 51185:4444
ip extcommunity-list standard CUSTOMER_EXTENDED_COMM_3016_IRBVRF3003_IPv6 permit rt 51185:3333
ip extcommunity-list standard CUSTOMER_EXTENDED_COMM_3017_4byte_IRBVRF3003_IPv4 permit rt 51185:65534
ip extcommunity-list standard CUSTOMER_EXTENDED_COMM_3017_4byte_IRBVRF3003_IPv6 permit rt 51185:65534
ip extcommunity-list standard CUSTOMER_EXT_COMM_ELAN1901 permit rt 1901:1901
!
ip as-path access-list ASPATH-IRBVRF3003-IN-V4 permit ^65030$
ip as-path access-list ASPATH-IRBVRF3003-IN-V6 permit ^65020$
!
line console 0
 exec-timeout 0
!
!
end

!
PE2-7001#
```

RR1

```
RR1-7036#sh run
!
```

```
! Software version: EC_AS5912-54X-OcNOS-SP-MPLS-7.0.0.129-Alpha 10/06/2025 17:41:16
!
! Last configuration change at 16:51:57 UTC Tue Oct 07 2025 by root
!
!
service password-encryption
!
logging console 5
logging monitor 5
logging cli
logging logfile ts_issue07 7
logging level nsm 3
logging level rip 5
logging level ripng 5
logging level ospf 5
logging level ospf6 5
logging level isis 5
logging level hostp 3
logging level ldp 5
logging level rsvp 5
logging level mrib 5
logging level pim 5
logging level auth 5
logging level mstp 5
logging level onm 5
logging level hsl 3
logging level oam 5
logging level vlog 5
logging level vrrp 5
logging level ndd 5
logging level rib 5
logging level bgp 4
logging level l2mrib 5
logging level lag 5
logging level sflow 5
logging level cml 3
logging level pserv 5
logging level cmm 4
logging level all 4
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
snmp-server enable traps pwdelete
snmp-server enable traps pw
snmp-server enable traps mpls
snmp-server enable traps snmp authentication
snmp-server enable traps ospf
snmp-server enable traps bgp
snmp-server enable traps isis
!
bgp extended-asn-cap
!
forwarding profile kaps profile-two
hardware-profile statistics ingress-acl enable
!
bfd interval 3 minrx 3 multiplier 3
!
key chain isis
 key-id 3
  key-string encrypted 0xc8a471564ac751dc
!
key chain BGP
 key-id 4
  key-string encrypted 0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157be7f6e35f32bf637
!
qos enable
!
```

```
mpls lsp-stitching
mpls ilm-ecmp ldp
mpls ftn-ecmp ldp
mpls label mode vpnv4 all-vrfs per-vrf
mpls label mode vpnv6 all-vrfs per-vrf
mpls label mode all-afs all-vrfs per-vrf
!
hostname RR1-7036
ip domain-lookup vrf management
ip name-server vrf management 10.12.3.23
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
ospf restart grace-period 2
ospf restart helper max-grace-period 2
feature ssh vrf management
aaa local authentication attempts max-fail 25
aaa local authentication unlock-timeout 1
aaa authentication login error-enable
snmp-server enable snmp vrf management
snmp-server view all .1 included vrf management
snmp-server community public vrf management
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
feature rsyslog
logging remote server 10.14.103.230 5 port 1514 vrf management
logging remote server 10.16.100.20 5 port 1514 vrf management
logging remote server 10.16.100.20 5 port 1514
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!
router-id 12.12.12.12
!
evpn mpls enable
!
ip vrf management
!
segment-routing
!
ip multicast-routing
!
ipv6 multicast-routing
!
ip prefix-list PL-BGPLU
 seq 5 permit 101.101.101.101/32
 seq 10 permit 201.201.201.201/32
 seq 15 permit 13.13.13.13/32
 seq 20 permit 10.137.76.17/32
!
ip prefix-list PL-EVPN
 seq 5 permit 1.0.2.0/24
 seq 10 permit 1.0.1.0/24
!
ipv6 prefix-list PFX-EVPNv6
 seq 5 permit 2000:1:1:1::/64
 seq 10 permit 2000:1:1:2::/64
!
router ldp
 router-id 12.12.12.12
 fast-reroute
```

© 2026 IP Infusion Inc. Proprietary

```
 pw-status-tlv
!
router rsvp
!
route-map RM-EXPORT-BGPLU permit 10
 match ip address prefix-list PL-BGPLU
 set metric 333
 set community 1:1
!
route-map RM-EXPORT-BGPLU permit 20
!
route-map RM-EXPORT-EVPN permit 10
 match ip address prefix-list PL-EVPN
 set originator-id 11.12.13.14
!
route-map RM-EXPORT-EVPN permit 20
 match ipv6 address prefix-list PFX-EVPNv6
 set aigp-metric 234
 set originator-id 33.33.33.33
!
route-map RM-EXPORT-EVPN permit 50
!
interface ce49
 description connected_to_pe2
 load-interval 30
 ip address 203.0.113.18/31
 mtu 9194
 label-switching
 link-debounce-time 2000 0
 ip ospf network point-to-point
 ip ospf authentication message-digest
 ip ospf authentication-key 0xfff87e79fdacd4e7
 ip ospf message-digest-key 3 md5 0x4c945d5d950eb831
 ipv6 ospf network point-to-point instance-id 0
 ipv6 router ospf area 0.0.0.0 tag 100 instance-id 0
 isis network point-to-point
 ip router isis 1
 ipv6 router isis 1
 isis authentication mode md5 level-1
 isis authentication mode md5 level-2
 isis authentication key-chain isis level-1
 isis authentication key-chain isis level-2
 enable-ldp ipv4
 mpls ldp-igp sync-delay 30
 enable-rsvp
 ip pim sparse-mode
 lldp-agent
  set lldp enable txrx
  set lldp chassis-id-tlv ip-address
  set lldp port-id-tlv if-name
  lldp tlv basic-mgmt system-name select
  lldp tlv basic-mgmt system-description select
  exit
 bfd interval 10 minrx 10 multiplier 3
!
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
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface loopback1
 ip address 12.12.12.12/32
 ipv6 address cafe:2012:12::12/128
 prefix-sid index 7 explicit-null n-flag-clear
 ipv6 router ospf area 0.0.0.0 instance-id 0
 ip router isis 1
 ipv6 router isis 1
 ip pim sparse-mode
!
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
```

© 2026 IP Infusion Inc. Proprietary

```
!
interface xe26
 description connected_to_pe1
 load-interval 30
 ip address 203.0.113.17/31
 ipv6 address 2003:0:113::17/64
 mtu 9194
 label-switching
 link-debounce-time 2000 0
 isis network point-to-point
 ip router isis 1
 ipv6 router isis 1
 isis authentication mode md5 level-1
 isis authentication mode md5 level-2
 isis authentication key-chain isis level-1
 isis authentication key-chain isis level-2
 enable-ldp ipv4
 mpls ldp-igp sync-delay 30
 enable-rsvp
 ip pim sparse-mode
 lldp-agent
  set lldp enable txrx
  set lldp chassis-id-tlv ip-address
  set lldp port-id-tlv if-name
  lldp tlv basic-mgmt system-name select
  lldp tlv basic-mgmt system-description select
  exit
 bfd interval 10 minrx 10 multiplier 3
!
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
 description connected_to_cisco_port-5
 load-interval 30
 ip address 203.0.113.21/31
 label-switching
 isis network point-to-point
 ip router isis 1
 enable-ldp ipv4
 lldp-agent
  set lldp enable txrx
  set lldp chassis-id-tlv ip-address
  set lldp port-id-tlv if-name
  lldp tlv basic-mgmt system-name select
  lldp tlv basic-mgmt system-description select
  exit
 bfd interval 10 minrx 10 multiplier 3
!
 exit
!
router ospf 100
 fast-reroute keep-all-paths
 bfd all-interfaces
 fast-reroute per-prefix remote-lfa area 0.0.0.0 tunnel mpls-ldp
 network 12.12.12.12/32 area 0.0.0.0
 network 203.0.113.18/31 area 0.0.0.0
!
router isis 1
 is-type level-1
 ignore-lsp-errors
 lsp-gen-interval 5
 max-lsp-lifetime 2000
 spf-interval-exp level-2 50 2000
 metric-style wide
 microloop-avoidance level-1
 microloop-avoidance max-fib 60 level-1
 mpls traffic-eng router-id 12.12.12.12
 mpls traffic-eng level-1
 capability cspf
 dynamic-hostname
 fast-reroute terminate-hold-on interval 100000
 fast-reroute per-prefix level-2 proto ipv4 all
 fast-reroute per-prefix remote-lfa level-2 proto ipv4 tunnel mpls-ldp
 fast-reroute ti-lfa level-2 proto ipv4
 bfd all-interfaces
 net 49.0001.0000.1102.00
 isis segment-routing global block 16000 23999
 segment-routing entropy-label
!
router isis ISIS-IGP-100
 is-type level-1
 authentication mode md5 level-1
 authentication key-chain isis level-1
 ignore-lsp-errors
 lsp-gen-interval 5
 max-lsp-lifetime 2000
 spf-interval-exp level-1 50 2000
 metric-style wide
 microloop-avoidance level-1
 microloop-avoidance max-fib 60 level-1
 mpls traffic-eng router-id 12.12.12.12
 mpls traffic-eng level-1
 capability cspf
 dynamic-hostname
 fast-reroute terminate-hold-on interval 100000
 fast-reroute per-prefix level-1 proto ipv4 all
```

© 2026 IP Infusion Inc. Proprietary

```
 fast-reroute ti-lfa level-1 proto ipv4
 bfd all-interfaces
 net 49.0001.0100.0000.1018.00
 passive-interface
!
router isis OCNOS-CISCO
 is-type level-1
 ignore-lsp-errors
 lsp-gen-interval 5
 max-lsp-lifetime 2000
 spf-interval-exp level-2 50 2000
 metric-style wide
 microloop-avoidance level-1
 microloop-avoidance max-fib 60 level-1
 mpls traffic-eng router-id 12.12.12.12
 mpls traffic-eng level-1
 capability cspf
 dynamic-hostname
 fast-reroute terminate-hold-on interval 100000
 fast-reroute per-prefix level-2 proto ipv4 all
 fast-reroute per-prefix remote-lfa level-2 proto ipv4 tunnel mpls-ldp
 fast-reroute ti-lfa level-2 proto ipv4
 bfd all-interfaces
 net 49.0001.0000.0026.00
!
router bgp 4200000001
 bgp router-id 12.12.12.12
 bgp auto-policy-soft-reset enable
 bgp cluster-id 4200000001
 bgp log-neighbor-changes
 no bgp inbound-route-filter
 allocate-label all
 neighbor PG-RR-PE1 peer-group
 neighbor PG-RR-PE1 remote-as 4200000001
 neighbor PG-RR-PE1 tcp-mss 1440
 neighbor PG-RR-PE1 update-source loopback1
 neighbor PG-RR-PE1 authentication-key
0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157be7f6e35f32bf637
 neighbor PG-RR-PE1 advertisement-interval 0
 neighbor PG-RR-PE1 fall-over bfd multihop
 neighbor PG-RR-PE2 peer-group
 neighbor PG-RR-PE2 remote-as 4200000001
 neighbor PG-RR-PE2 tcp-mss 1440
 neighbor PG-RR-PE2 update-source loopback1
 neighbor PG-RR-PE2 authentication-key
0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157be7f6e35f32bf637
 neighbor PG-RR-PE2 advertisement-interval 0
 neighbor PG-RR-PE2 fall-over bfd multihop
 neighbor 13.13.13.13 remote-as 65002
 neighbor 13.13.13.13 ebgp-multihop 255
 neighbor 13.13.13.13 update-source loopback1
 neighbor 101.101.101.101 peer-group PG-RR-PE1
 neighbor 201.201.201.201 peer-group PG-RR-PE2
 !
 address-family ipv4 unicast
 network 12.12.0.12/32
 network 12.12.12.12/32
 exit-address-family
 !
 address-family ipv4 labeled-unicast
 neighbor PG-RR-PE1 activate
 neighbor PG-RR-PE1 route-reflector-client
 neighbor PG-RR-PE1 route-map RM-EXPORT-BGPLU out
 neighbor PG-RR-PE2 activate
 neighbor PG-RR-PE2 route-reflector-client
 neighbor PG-RR-PE2 route-map RM-EXPORT-BGPLU out
 neighbor 13.13.13.13 activate
 exit-address-family
```

© 2026 IP Infusion Inc. Proprietary

```
 !
 address-family vpnv4 unicast
 neighbor PG-RR-PE1 activate
 neighbor PG-RR-PE1 route-reflector-client
 neighbor PG-RR-PE1 next-hop-self
 neighbor PG-RR-PE2 activate
 neighbor PG-RR-PE2 route-reflector-client
 neighbor PG-RR-PE2 next-hop-self
 neighbor PG-RR-PE2 route-map RM-EXPORT-EVPN out
 neighbor 13.13.13.13 allow-ebgp-vpn
 neighbor 13.13.13.13 activate
 neighbor 13.13.13.13 aigp enable
 neighbor 13.13.13.13 route-map RM-EXPORT-EVPN out
 exit-address-family
 !
 address-family rtfilter unicast
 exit-address-family
 !
 address-family l2vpn vpls
 neighbor PG-RR-PE1 activate
 neighbor PG-RR-PE1 route-reflector-client
 neighbor PG-RR-PE2 activate
 neighbor PG-RR-PE2 route-reflector-client
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor PG-RR-PE1 activate
 neighbor PG-RR-PE1 route-reflector-client
 neighbor PG-RR-PE2 activate
 neighbor PG-RR-PE2 route-reflector-client
 exit-address-family
 !
 address-family vpnv6 unicast
 neighbor PG-RR-PE1 activate
 neighbor PG-RR-PE1 route-reflector-client
 neighbor PG-RR-PE2 activate
 neighbor PG-RR-PE2 route-reflector-client
 neighbor PG-RR-PE2 route-map RM-EXPORT-EVPN out
 exit-address-family
 !
 address-family ipv6 unicast
 redistribute connected
 exit-address-family
 !
 address-family ipv6 labeled-unicast
 neighbor PG-RR-PE1 activate
 neighbor PG-RR-PE1 route-reflector-client
 neighbor PG-RR-PE2 activate
 neighbor PG-RR-PE2 route-reflector-client
 exit-address-family
 !
 exit
!
!
end

!
RR1-7036#
```

RR1

```
RR#sh run
!
! Software version: EC_AS5912-54X-OcNOS-SP-MPLS-7.0.0.261-GA 02/20/2026 15:31:03
!
! Last configuration change at 17:14:02 UTC Tue Feb 24 2026 by root
!
```

```
service password-encryption
!
logging console 3
logging monitor 5
logging logfile device_debug_log 2
logging level nsm 5
logging level ospf 5
logging level ldp 5
logging level rsvp 5
logging level hsl 5
logging level bgp 5
logging level cml 5
logging level all 5
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
bgp extended-asn-cap
!
forwarding profile kaps profile-two
hardware-profile filter qos enable
hardware-profile statistics ingress-acl enable
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
hostname RR
tfo Disable
errdisable cause stp-bpdu-guard
enable ext-ospf-multi-inst
feature dns relay
ip dns relay
ipv6 dns relay
!
ip vrf management
!
ip prefix-list PL-BGPLU
 seq 5 permit 1.1.1.1/32
 seq 10 permit 3.3.3.3/32
!
router ldp
 router-id 2.2.2.2
!
router rsvp
!
route-map RM-EXPORT-BGPLU permit 10
 match ip address prefix-list PL-BGPLU
 set ip next-hop self
!
route-map RM-EXPORT-BGPLU permit 20
!
interface ce49
 load-interval 30
 ip address 30.1.1.2/24
 mtu 9216
 label-switching
 ip ospf network point-to-point
 enable-rsvp
!
interface ce50
!
interface ce51
!
interface ce52
!
interface ce53
```

© 2026 IP Infusion Inc. Proprietary

```
!
interface ce54
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ip address 2.2.2.2/32 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
!
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
```

```
 load-interval 30
 ip address 10.1.1.2/24
 mtu 9216
 label-switching
 ip ospf network point-to-point
 enable-ldp ipv4
!
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
!
router ospf 65530
 ospf router-id 2.2.2.2
 bfd all-interfaces
 network 2.2.2.2/32 area 0.0.0.0 instance-id 1
 network 30.1.1.0/24 area 0.0.0.0
!
router ospf 65535
 ospf router-id 2.2.2.2
 bfd all-interfaces
 network 2.2.2.2/32 area 0.0.0.0
 network 10.1.1.0/24 area 0.0.0.0
!
router bgp 4200000001
 bgp router-id 2.2.2.2
 bgp auto-policy-soft-reset enable
```

© 2026 IP Infusion Inc. Proprietary

```
 bgp cluster-id 2.2.2.2
 bgp log-neighbor-changes
 no bgp inbound-route-filter
 allocate-label all
 neighbor 1.1.1.1 remote-as 4200000001
 neighbor 1.1.1.1 tcp-mss 1440
 neighbor 1.1.1.1 update-source 2.2.2.2
 neighbor 1.1.1.1 authentication-key
0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157be7f6e35f32bf637
 neighbor 1.1.1.1 advertisement-interval 0
 neighbor 1.1.1.1 fall-over bfd multihop
 neighbor 3.3.3.3 remote-as 4200000001
 neighbor 3.3.3.3 tcp-mss 1440
 neighbor 3.3.3.3 update-source 2.2.2.2
 neighbor 3.3.3.3 authentication-key
0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157be7f6e35f32bf637
 neighbor 3.3.3.3 advertisement-interval 0
 neighbor 3.3.3.3 fall-over bfd multihop
 !
 address-family ipv4 unicast
 network 2.2.2.2/32
 exit-address-family
 !
 address-family ipv4 labeled-unicast
 neighbor 1.1.1.1 activate
 neighbor 1.1.1.1 route-reflector-client
 neighbor 1.1.1.1 route-map RM-EXPORT-BGPLU out
 neighbor 3.3.3.3 activate
 neighbor 3.3.3.3 route-reflector-client
 neighbor 3.3.3.3 route-map RM-EXPORT-BGPLU out
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 1.1.1.1 activate
 neighbor 1.1.1.1 route-reflector-client
 neighbor 3.3.3.3 activate
 neighbor 3.3.3.3 route-reflector-client
 exit-address-family
 !
 exit
!
rsvp-trunk P1_to_PE3 ipv4
 to 3.3.3.3
!
line console 0
 exec-timeout 0 0
line vty 0 16
 exec-timeout 0 0
!
!
end
```

Validation

```
Establish end-to-end PE1-PE2 labeled reachability using BGP-LU across dual OSPF domains (PE1-RR with
OSPF+LDP and RR-PE2 with OSPF+RSVP)


## Verification


## ! Define prefix-lists on RR node
RR#sh run prefix-list
!
ip prefix-list PL-BGPLU
 seq 5 permit 1.1.1.1/32
 seq 10 permit 3.3.3.3/32
```

```
!



## ! Define route-map for labeled-unicast

RR#sh run route-map
!
route-map RM-EXPORT-BGPLU permit 10
 match ip address prefix-list PL-BGPLU
 set ip next-hop self
!
route-map RM-EXPORT-BGPLU permit 20
!


# Verify prefix-list and route-map
RR#sh ip prefix-list detail PL-BGPLU
ip prefix-list PL-BGPLU:
   count: 2, range entries: 0, sequences: 5 - 10
    ripd:
     seq 5 permit 1.1.1.1/32 (hit count: 0, refcount: 0)
     seq 10 permit 3.3.3.3/32 (hit count: 0, refcount: 0)
    ripngd:
     seq 5 permit 1.1.1.1/32 (hit count: 0, refcount: 0)
     seq 10 permit 3.3.3.3/32 (hit count: 0, refcount: 0)
    ospfd:
     seq 5 permit 1.1.1.1/32 (hit count: 0, refcount: 0)
     seq 10 permit 3.3.3.3/32 (hit count: 0, refcount: 0)
    ospf6d:
     seq 5 permit 1.1.1.1/32 (hit count: 0, refcount: 0)
     seq 10 permit 3.3.3.3/32 (hit count: 0, refcount: 0)
    ldpd:
     seq 5 permit 1.1.1.1/32 (hit count: 0, refcount: 0)
     seq 10 permit 3.3.3.3/32 (hit count: 0, refcount: 0)
    bgpd:
     seq 5 permit 1.1.1.1/32 (hit count: 1, refcount: 1)
     seq 10 permit 3.3.3.3/32 (hit count: 1, refcount: 1)
RR#sh route
route-map  router-id
RR#sh route-map
RR#sh route-map RM-EXPORT-BGPLU
route-map RM-EXPORT-BGPLU, permit, sequence 10
  Match clauses:
    ip address prefix-list: PL-BGPLU
  Set clauses:
    ip next-hop self
route-map RM-EXPORT-BGPLU, permit, sequence 20
  Match clauses:
  Set clauses:




## Verify on RR for next hop changed to itself (RR)
RR#sh ip bgp neighbors 1.1.1.1 advertised-routes

For address family: IPv4 Labeled-Unicast vrf: default
BGP table version is 3, local router ID is 2.2.2.2
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
Description : Ext-Color - Extended community color

    Network          Next Hop            Metric    LocPrf   Weight Path   Ext-Color
*>i   2.2.2.2/32       2.2.2.2              0        100       32768  i        -
```

```
*>i   3.3.3.3/32       2.2.2.2              0        100       0    i        -
Total number of prefixes 2


RR#sh ip bgp labeled-unicast  summary
BGP router identifier 2.2.2.2, local AS number 4200000001
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   Desc
1.1.1.1            4 4200000001         34         39       3      0      0  00:12:04            1
3.3.3.3            4 4200000001         35         43       3      0      0  00:12:33            1

Total number of neighbors 2

Total number of Established sessions 2

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

RR#sh ip bgp neighbors 3.3.3.3 advertised-routes

For address family: IPv4 Labeled-Unicast vrf: default
BGP table version is 3, local router ID is 2.2.2.2
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
Description : Ext-Color - Extended community color

    Network          Next Hop            Metric    LocPrf   Weight Path   Ext-Color
*>i   1.1.1.1/32       2.2.2.2              0        100       0    i        -
*>i   2.2.2.2/32       2.2.2.2              0        100       32768  i        -
Total number of prefixes 2



RR#sh mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled
Code    FEC                 FTN-ID    Nhlfe-ID  Tunnel-ID   Pri   Out-Label    Out-
Intf    ELC       Nexthop         Algo-Num   UpTime
   L>   1.1.1.1/32          3         4         -           -     -            -           -
         -               N/A        00:19:32
                                      3         -
           Yes   3            xe26        No        10.1.1.1        -          -
   B    1.1.1.1/32          4         5         -           Yes   24324        -
           No        1.1.1.1         N/A        -
R
(t)>   3.3.3.3/32          1         1         5001        Yes   24320        ce49        No        3
0.1.1.1        N/A        00:20:01
   B    3.3.3.3/32          2         2         -           Yes   24962        -
           No        3.3.3.3         N/A        -
RR#sh mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP , K - CLI , V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy,       U - unknown, UPStr - upstream

ILM-ECMP LDP: Disabled, SR: Disabled
Code    FEC/VRF/L2CKT    ILM-ID      In-Label    Out-Label   In-Intf    Out-
Intf/VRF       Nexthop                 pri  Algo-Num  UpTime    UPStr peers
```

```
   B>   2.2.2.2/32         2           26240       Nolabel     N/A        N/A              127.0.0.1
              Yes  N/A       00:19:57
   R>   2.2.2.2/32         1           24320       Nolabel     N/A        N/A              127.0.0.1
              Yes  N/A       00:19:59  1
   L>   3.3.3.3/32         4           24961       Nolabel     N/A        N/A              127.0.0.1
              Yes  N/A       00:19:36  1
   B>   3.3.3.3/32         3           26241       24962       N/A        N/A              3.3.3.3
              Yes  N/A       00:19:54
   B>   1.1.1.1/32         5           26242       24324       N/A        N/A              1.1.1.1
              Yes  N/A       00:19:25
RR#




================================

## Verification on PE1 Next hop changed to RR Loopback address

PE1#sh ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal, S - stale
     Network          Next Hop        In Label     Out Label
*>  1.1.1.1/32       0.0.0.0           24324          -
*>i 2.2.2.2/32       2.2.2.2           24323        26240
*>i 3.3.3.3/32       2.2.2.2           24322        26241
PE1#sh ip bgp labeled-unicast summary
BGP router identifier 1.1.1.1, local AS number 4200000001
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   Desc
2.2.2.2            4 4200000001         43         38       2      0      0  00:13:41            2

Total number of neighbors 1

Total number of Established sessions 1

PE1#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            1.1.1.1/32 is directly connected, lo, installed 00:14:39, last update 00:14:39 ago
O            2.2.2.2/32 [110/2] via 10.1.1.2, xe26, installed 00:14:25, last update 00:14:25 ago
B            3.3.3.3/32 [200/0] via 2.2.2.2 (recursive via 10.1.1.2), installed 00:04:33, last update
00:04:33 ago
C            10.1.1.0/24 is directly connected, xe26, installed 00:14:40, last update 00:14:40 ago
C            127.0.0.0/8 is directly connected, lo, installed 00:22:06, last update 00:22:06 ago

Gateway of last resort is not set


PE1#sh mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP
```

```
FTN-ECMP LDP: Disabled, SR: Disabled
Code    FEC                 FTN-ID    Nhlfe-ID  Tunnel-ID   Pri   Out-Label    Out-
Intf    ELC       Nexthop         Algo-Num   UpTime
   L>   2.2.2.2/32          3         6         -           -     -            -           -
         -               N/A        00:20:03
                                      5         -
           Yes   3           xe26        No        10.1.1.2        -          -
   B    2.2.2.2/32          2         4         -           Yes   26240        -
           No        2.2.2.2         N/A       -
   B>   3.3.3.3/32          1         12        -
           Yes   26241       xe26        No        2.2.2.2         N/A        00:10:31
PE1#sh mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP , K - CLI , V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy,       U - unknown, UPStr - upstream

ILM-ECMP LDP: Disabled, SR: Disabled
Code    FEC/VRF/L2CKT    ILM-ID      In-Label    Out-Label   In-Intf    Out-
Intf/VRF       Nexthop                pri  Algo-Num  UpTime    UPStr peers
   B>   VRF2               2           24321       Nolabel     N/A        N/A              N/A
            Yes  N/A       00:20:13
   B>   VRF1               1           24320       Nolabel     N/A        N/A              N/A
            Yes  N/A       00:20:13
   B>   2.2.2.2/32         4           24323       26240       N/A        N/A              2.2.2.2
            Yes  N/A       00:20:06
   B>   3.3.3.3/32         3           24322       26241       N/A        N/A              2.2.2.2
            Yes  N/A       00:10:34
   B>   1.1.1.1/32         5           24324       Nolabel     N/A        N/A              127.0.0.1
            Yes  N/A       00:20:09
PE1#ping ip 3.3.3.3
Press CTRL+C to exit
PING 3.3.3.3 (3.3.3.3) 100(128) bytes of data.
108 bytes from 3.3.3.3: icmp_seq=1 ttl=64 time=0.641 ms
108 bytes from 3.3.3.3: icmp_seq=2 ttl=64 time=0.537 ms
108 bytes from 3.3.3.3: icmp_seq=3 ttl=64 time=0.669 ms
108 bytes from 3.3.3.3: icmp_seq=4 ttl=64 time=0.546 ms
108 bytes from 3.3.3.3: icmp_seq=5 ttl=64 time=0.606 ms

--- 3.3.3.3 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4132ms
rtt min/avg/max/mdev = 0.537/0.599/0.669/0.051 ms




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


## Verify on PE2

PE2#sh ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.1/32       2.2.2.2          24964        26242
*>i 2.2.2.2/32       2.2.2.2          24963        26240
*>  3.3.3.3/32       0.0.0.0          24962          -
PE2#sh ip bgp labeled-unicast summary
BGP router identifier 3.3.3.3, local AS number 4200000001
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   Desc
2.2.2.2            4 4200000001         47         40       3      0      0  00:14:38            2
```

```
Total number of neighbors 1

Total number of Established sessions 1
PE2#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            1.1.1.1/32 [200/0] via 2.2.2.2 (recursive via 30.1.1.2), installed 00:03:59, last update
00:03:59 ago
O            2.2.2.2/32 [110/2] via 30.1.1.2, ce0, installed 00:14:44, last update 00:14:44 ago
C            3.3.3.3/32 is directly connected, lo, installed 00:14:59, last update 00:14:59 ago
C            30.1.1.0/24 is directly connected, ce0, installed 00:14:59, last update 00:14:59 ago
C            127.0.0.0/8 is directly connected, lo, installed 00:21:32, last update 00:21:32 ago

Gateway of last resort is not set


PE2#sh mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled
Code    FEC                 FTN-ID    Nhlfe-ID  Tunnel-ID   Pri   Out-Label    Out-
Intf    ELC       Nexthop         Algo-Num   UpTime
   B>   1.1.1.1/32          3         11        -
          Yes   26242        ce0         No        2.2.2.2         N/A        00:11:11
R
(t)>   2.2.2.2/32          1         3         5001        Yes   24320        ce0         No        3
0.1.1.2        N/A        00:21:09
   B    2.2.2.2/32           2         4         -           Yes   26240        -
          No        2.2.2.2         N/A        -
PE2#sh mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP , K - CLI , V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy,       U - unknown, UPStr - upstream

ILM-ECMP LDP: Disabled, SR: Disabled
Code    FEC/VRF/L2CKT    ILM-ID      In-Label    Out-Label   In-Intf    Out-
Intf/VRF       Nexthop                 pri  Algo-Num  UpTime    UPStr peers
   B>   3.3.3.3/32         4           24962       Nolabel     N/A        N/A              127.0.0.1
          Yes  N/A       00:21:12
   B>   VRF1               1           24960       Nolabel     N/A        N/A              N/A
          Yes  N/A       00:21:22
   R>   3.3.3.3/32         3           24320       Nolabel     N/A        N/A              127.0.0.1
          Yes  N/A       00:21:18  1
   B>   VRF2               2           24961       Nolabel     N/A        N/A              N/A
          Yes  N/A       00:21:22
   B>   2.2.2.2/32         5           24963       26240       N/A        N/A              2.2.2.2
          Yes  N/A       00:21:06
   B>   1.1.1.1/32         6           24964       26242       N/A        N/A              2.2.2.2
          Yes  N/A       00:11:14
PE2#ping ip 1.1.1.1
Press CTRL+C to exit
PING 1.1.1.1 (1.1.1.1) 100(128) bytes of data.
108 bytes from 1.1.1.1: icmp_seq=1 ttl=64 time=0.672 ms
```

```
108 bytes from 1.1.1.1: icmp_seq=2 ttl=64 time=0.491 ms
108 bytes from 1.1.1.1: icmp_seq=3 ttl=64 time=0.516 ms
108 bytes from 1.1.1.1: icmp_seq=4 ttl=64 time=0.515 ms
108 bytes from 1.1.1.1: icmp_seq=5 ttl=64 time=0.516 ms

--- 1.1.1.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4103ms
rtt min/avg/max/mdev = 0.491/0.542/0.672/0.065 ms
```

Commands

The BGP LU Next-hop self in Route-map feature has the following configuration commands:

set ip next-hop self

Use this command to set next hop self for IPV4 BGP-LU neighbors.

Use “no” form of this command to not set next hop self.

Command Syntax

```
set ip next-hop self
no set ip next-hop self
```

Parameters

None

Command Mode

Route map mode mode

Applicability

This command is introduced in OcNOS version 7.0.0

Examples

```
ip prefix-list BCOM-IP
   seq 5 permit 25.4.4.0/24 eq 24

route-map BCOM-RM permit 2
   match ip address prefix-list BCOM-IP
   set ip next-hop self



router bgp 26
..
  address-family ipv4 labeled-unicast
  neighbor 27.27.27.27 route-map BCOM-RM out
  ..
```

BGP Additional Paths for VPNv4

Overview

The Border Gateway Protocol (BGP) ADDPATH feature allows the advertisement of multiple paths through the same peering 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.

Normal BGP Behavior

By default, all BGP routers and Route-Reflectors propagate only their best paths over their sessions. In case they advertise any route with the same NLRI as a previously advertised route, the latest one implicitly replaces the previous advertisement, which is known as an Implicit Withdraw. The Implicit Withdraw can achieve better scaling, but at the cost of path diversity.

The use of route-reflectors (or confederations), thus has significant effect on redundancy by hiding alternate paths. Using full-mesh is not an option, so a mechanism is needed to allow the propagation of multiple alternate paths in an RR/Confederation environment. Such mechanism is already available in BGP/MPLS VPN scenarios, where multiple point of attachments for CE sites could utilize different RD values to differentiate the same routes advertised from different connection points. However, a generic solution is required, allowing for advertising multiple alternate paths with IPv4 or any other address-family.

The “Advertisement of Multiple Paths in BGP” or “BGP Add-Path” as the feature is usually called is a BGP extension that allows the advertisement of multiple paths for the same address prefix without the new paths implicitly replacing any previously advertised ones.

BGP Behavior with ADDPATH

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.

Address Family Identifier (AFI)	2 octets
Subsequent Address Family Identifier (SAFI)	1 octet
Send/Receive	1 octet

The send/receive field in the BGP Capability TLV indicates whether for a given <AFI, SAFI>, the sender is able to :

  - Receive multiple paths from its peer (value 1)
  - Send multiple paths to its peer (value 2), or
  - both (value 3)
- Each alternate path is identified by a Path Identifier in addition to the address prefix

Path Identifier	4 octets
Length	1 octet
Prefix	variable

© 2026 IP Infusion Inc. Proprietary

In the event of a next-hop failure, the BGP Add-Path feature hence improves the BGP control plane convergence.

Notes:

- BGP does not advertise additional (non-best) paths to an eBGP neighbor, or to an iBGP neighbor when **`next-hop-self`** is explicitly configured (using the **`neighbor A.B.C.D next-hop-self`** command). Furthermore, if the best path is learned from an eBGP neighbor, any other internal paths if available for the same routes will not be selected or considered as additional paths.
- We neither support BGP Add-path for eBGP neighbors, nor do we currently have a use case for it. Although our documentation previously included eBGP Add-path, the feature primarily applies to iBGP Route Reflectors (RRs) with in VPN environments.

Topology

Figure 83. BGP Add-Path VPNv4 Topology

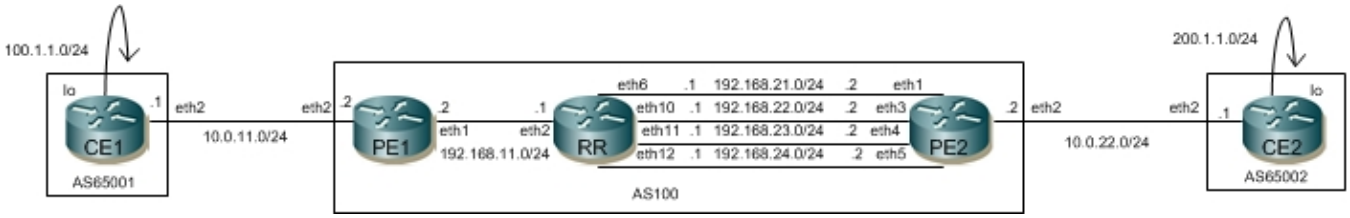

Initial Configuration

CE1 Configuration

`#configure terminal`	Enter configure mode
`(config)#interface eth2`	Enter Interface mode for eth2.
`(config-if)#ip address 10.0.11.1/24`	Assign IP address.
`(config-if)#exit`	Exit interface mode.
`(config)#interface lo`	Enter Interface mode for lo.
`(config-if)#ip address 100.1.1.2/24`	Assign IP address.
`(config-if)#exit`	Exit interface mode.
`(config)#router bgp 65001`	Enter BGP router mode.
`(config-router)#neighbor 10.0.11.2 remote-as 100`	Configure an eBGP neighbor on PE1
`(config-router)#network 100.1.1.0/24`	Announce the network 100.1.1.0/24 into BGP
`(config-router)#address-family ipv4 unicast`	Configure the address family ipv4 under router BGP
`(config-router-af)#neighbor 10.0.11.2 activate`	Activate eBGP neighbor on PE1
`(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.

(config-router)#exit	Exit Router BGP.
(config)#commit	Apply commit
(config)#exit	Exit Router BGP and Configure mode.

PE1Configuration

#configure terminal	Enter configure mode
(config)#ip vrf vrf1	Configure a VRF vrf1
(config-vrf)#rd 100:1	Configure the Route-Distinguisher ie., RD for vrf1
(config-vrf)#route-target both 200:1	Configure the Route-Target ie., RT to import and export the VPNv4 routes
(config-vrf)#exit	Exit the configure VRF mode
(config)#interface eth2	Configure the interface eth2
(config-if)#ip vrf forwarding vrf1	Configure the interface eth2 for IP VRF forwarding for vrf1
(config-if)#ip address 10.0.11.2/24	Configure the IP address for interface eth2
(config-if)#exit	Exit the interface mode for interface eth2
(config)#interface eth1	Configure the interface eth1
(config-if)#ip address 192.168.11.2/24	Configure the IP address for interface eth1
(config-if)#exit	Exit the interface mode for eth1
(config)#router bgp 100	Configure the BGP routing instance 100
(config-router)#neighbor 192.168.11.1 remote-as 100	Configure the neighbor address and remote-as for the 4 iBGP neighbors on RR
(config-router)#address-family vpnv4 unicast	Configure the address family VPNv4 under router BGP
(config-router-af)#neighbor 192.168.11.1 activate	Activate the IPv4 iBGP neighbors on RR for the VPNv4 address family
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#address-family ipv4 vrf vrf1	Configure the address family IPv4 VRF vrf1 under router BGP
(config-router-af)#neighbor 10.0.11.1 remote-as 65002	Configure the neighbor address and remote-as for the eBGP neighbor on CE1
(config-router-af)#neighbor 10.0.11.1 activate	Activate the eBGP neighbor on CE1 for the address family IPv4 VRF vrf1
(config-router-af)#exit	Exit the router BGP mode and return to the configure terminal mode
(config)#router ospf 100	Configure an OSPF routing instance
(config-router)#network 192.168.11.0/24 area 0	Define the interface eth1 to run OSPF and associate the area ID 0 with interface eth1

© 2026 IP Infusion Inc. Proprietary

(config-router)#exit	Exit the router OSPF mode and return to the configure terminal mode
(config)#router ldp	
(config-router)#exit	Exit the Router mode for LDP and return to the configure terminal mode
(config)#interface eth1	Enter interface mode for interfacec eth1
(config-if)#label-switching	Enable label-switching for interface eth1
(config-if)#enable-ldp ipv4	Enable LDP for IPv4 on interface eth1
(config)#commit	Apply commit
(config)#exit	Exit Router BGP and Configure mode.

RR Configuration

#configure terminal	Enter configure mode
(config)#interface eth2	Enter interface mode for interface eth2
(config-if)#ip address 192.168.11.1/24	Configure IP address for the interface eth2
(config-if)#exit	Exit the interface mode for eth2
(config)#interface eth6	Enter interface mode for interface eth6
(config-if)#ip address 192.168.21.1/24	Configure IP address for the interface eth6
(config-if)#exit	Exit the interface mode for eth6
(config)#interface eth10	Enter interface mode for interface eth10
(config-if)#ip address 192.168.22.1/24	Configure IP address for the interface eth10
(config-if)#exit	Exit the interface mode for eth10
(config)#interface eth11	Enter interface mode for interface eth11
(config-if)#ip address 192.168.23.1/24	Configure IP address for the interface eth11
(config-if)#exit	Exit the interface mode for eth11
(config)#interface eth12	Enter interface mode for interface eth12
(config-if)#ip address 192.168.24.1/24	Configure IP address for the interface eth12
(config-if)#exit	Exit the interface mode for eth12
(config)#router bgp 100	Enter the router BGP mode
(config-router)#neighbor 192.168.11.2 remote-as 100	Configure the neighbor address and remote-as for the iBGP neighbor on PE1
(config-router)#neighbor 192.168.21.2 remote-as 100	Configure the neighbor address and remote-as for the iBGP neighbor on interface eth1 of PE2
(config-router)#neighbor 192.168.22.2 remote-as 100	Configure the neighbor address and remote-as for the iBGP neighbor on interface eth3 of PE2
(config-router)#neighbor 192.168.23.2 remote-as 100	Configure the neighbor address and remote-as for the iBGP neighbor on interface eth4 of PE2
(config-router)#neighbor 192.168.24.2 remote-as 100	Configure the neighbor address and remote-as for

	the iBGP neighbor on interface eth5 of PE2
`(config-router)#address-family vpnv4 unicast`	Configure the address family VPNv4 under router BGP
`(config-router-af)#neighbor 192.168.11.2 activate`	Activate the IPv4 iBGP neighbor on PE1 for the VPNv4 address family
`(config-router-af)#neighbor 192.168.21.2 activate`	Activate the IPv4 iBGP neighbor on interface eth1 of PE2 for the VPNv4 address family
`(config-router-af)#neighbor 192.168.22.2 activate`	Activate the IPv4 iBGP neighbor on interface eth3 of PE2 for the VPNv4 address family
`(config-router-af)#neighbor 192.168.23.2 activate`	Activate the IPv4 iBGP neighbor on interface eth4 of PE2 for the VPNv4 address family
`(config-router-af)#neighbor 192.168.24.2 activate`	Activate the IPv4 iBGP neighbor on interface eth5 of PE2 for the VPNv4 address family
`(config-router-af)#neighbor 192.168.11.2 route-reflector-client`	Configure the VPNv4 neighbor on PE1 as route-reflector-client
`(config-router-af)#neighbor 192.168.21.2 route-reflector-client`	Configure the VPNv4 neighbor on interface eth1 of PE2 as route-reflector-client
`(config-router-af)#neighbor 192.168.22.2 route-reflector-client`	
	Configure the VPNv4 neighbor on interface eth3 of PE2 as route-reflector-client
`(config-router-af)#neighbor 192.168.23.2 route-reflector-client`	Configure the VPNv4 neighbor on interface eth4 of PE2 as route-reflector-client
`(config-router-af)#neighbor 192.168.24.2 route-reflector-client`	Configure the VPNv4 neighbor on interface eth5 of PE2 as route-reflector-client
`(config-router-af)#exit`	Exit the router BGP mode
`(config)#router ospf 100`	Configure an OSPF instance 100
`(config-router)#network 192.168.11.0/24 area 0`	Define the interface eth2 connected to PE1 to run OSPF and associate it with the backbone area ID 0
`(config-router)#network 192.168.21.0/24 area 0`	Define the interface eth6 connected to PE2 to run OSPF and associate it with the backbone area ID 0
`(config-router)#network 192.168.22.0/24 area 0`	Define the interface eth10 connected to PE2 to run OSPF and associate it with the backbone area ID 0
`(config-router)#network 192.168.23.0/24 area 0`	Define the interface eth11 connected to PE2 to run OSPF and associate it with the backbone area ID 0
`(config-router)#network 192.168.24.0/24 area 0`	Define the interface eth12 connected to PE2 to run OSPF and associate it with the backbone area ID 0
`(config-router)#exit`	Exit the router OSPF mode
`(config)#router ldp`	
`(config-router)#exit`	Exit the Router mode for LDP
`(config)#interface eth1`	Enter interface mode for interface eth2

© 2026 IP Infusion Inc. Proprietary

(config-if)#label-switching	Enable label-switching for interface eth2
(config-if)#enable-ldp ipv4	Enable LDP for IPv4 on interface eth2
(config-if)#exit	Exit the interface mode for eth2
(config)#interface eth6	Enter interface mode for interface eth6
(config-if)#label-switching	Enable label-switching for interface eth6
(config-if)#enable-ldp ipv4	Enable LDP for IPv4 on interface eth6
(config-if)#exit	Exit the interface mode for eth6
(config)#interface eth10	Enter interface mode for interface eth10
(config-if)#label-switching	Enable label-switching for interface eth10
(config-if)#enable-ldp ipv4	Enable LDP for IPv4 on interface eth10
(config-if)#exit	Exit the interface mode for eth10
(config)#interface eth11	Enter interface mode for interface eth11
(config-if)#label-switching	Enable label-switching for interface eth11
(config-if)#enable-ldp ipv4	Enable LDP for IPv4 on interface eth11
(config-if)#exit	Exit the interface mode for eth11
(config)#interface eth12	Enter interface mode for interface eth12
(config-if)#label-switching	Enable label-switching for interface eth12
(config-if)#enable-ldp ipv4	Enable LDP for IPv4 on interface eth12
(config)#commit	Apply commit
(config)#exit	exit the global mode

PE2 Configuration

#configure terminal	Enter configure mode
(config)#ip vrf vrf1	Configure a VRF vrf1
(config-vrf)#rd 100:1	Configure the Route-Distinguisher ie., RD for vrf1
(config-vrf)#route-target both 200:1	Configure the Route-Target ie., RT to import and export the VPNv4 routes
(config-vrf)#exit	Exit the configure VRF mode
(config)#interface eth2	Configure the interface eth2
(config-if)#ip vrf forwarding vrf1	Configure the interface eth2 for IP VRF forwarding for vrf1
(config-if)#ip address 10.0.22.2/24	Configure the IP address for interface eth2
(config-if)#exit	Exit the interface mode for eth2
(config)#interface eth1	Configure the interface eth1
(config-if)#ip address 192.168.21.2/24	Configure the IP address for interface eth1
(config-if)#exit	Exit the interface mode for eth1
(config)#interface eth3	Configure the interface eth3

© 2026 IP Infusion Inc. Proprietary

`(config-if)#ip address 192.168.22.2/24`	Configure the IP address for interface eth3
`(config-if)#exit`	Exit the interface mode for eth3
`(config)#interface eth4`	Configure the interface eth4
`(config-if)#ip address 192.168.23.2/24`	Configure the IP address for interface eth4
`(config-if)#exit`	Exit the interface mode for eth4
`(config)#interface eth5`	Configure the interface eth5
`(config-if)#ip address 192.168.24.2/24`	Configure the IP address for interface eth5
`(config-if)#exit`	Exit the interface mode for eth5
`(config)#router bgp 100`	Configure the BGP routing instance 100
`(config-router)#neighbor 192.168.21.1 remote-as 100`	Configure the neighbor address and remote-as for the iBGP neighbor on interface eth6 of PE2
`(config-router)#neighbor 192.168.22.1 remote-as 100`	Configure the neighbor address and remote-as for the iBGP neighbor on interface eth10 of PE2
`(config-router)#neighbor 192.168.23.1 remote-as 100`	Configure the neighbor address and remote-as for the iBGP neighbor on interface eth11 of PE2
`(config-router)#neighbor 192.168.24.1 remote-as 100`	Configure the neighbor address and remote-as for the iBGP neighbor on interface eth12 of PE2
`(config-router)#address-family vpnv4 unicast`	Configure the address family VPNv4 under router BGP
`(config-router-af)#neighbor 192.168.21.1 activate`	Activate the IPv4 iBGP neighbor on interface eth6 of PE2 for the VPNv4 address family
`(config-router-af)#neighbor 192.168.22.1 activate`	Activate the IPv4 iBGP neighbor on interface eth10 of PE2 for the VPNv4 address family
`(config-router-af)#neighbor 192.168.23.1 activate`	Activate the IPv4 iBGP neighbor on interface eth11 of PE2 for the VPNv4 address family
`(config-router-af)#neighbor 192.168.24.1 activate`	Activate the IPv4 iBGP neighbor on interface eth12 of PE2 for the VPNv4 address family
`(config-router-af)#exit`	Exit the router BGP address family VPNv4 mode
`(config)#router bgp 100`	Enter the router BGP mode
`(config-router)#address-family ipv4 vrf vrf1`	Configure the address family IPv4 VRF vrf1 under router BGP
`(config-router-af)#neighbor 10.0.22.1 remote-as 65002`	Configure the neighbor address and remote-as for the eBGP neighbor on CE2
`(config-router-af)#neighbor 10.0.22.1 activate`	Activate the 4 eBGP neighbor on CE2 for the address family IPv4 VRF vrf1
`(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`(config)#router ospf 100`	Configure an OSPF instance 100
`(config-router)#network 192.168.21.0/24 area 0`	Define the interface eth1 connected to PE2 to run OSPF and associate it with the backbone area ID 0

© 2026 IP Infusion Inc. Proprietary

`(config-router)#network 192.168.22.0/24 area 0`	Define the interface eth3 connected to PE2 to run OSPF and associate it with the backbone area ID 0
`(config-router)#network 192.168.23.0/24 area 0`	Define the interface eth4 connected to PE2 to run OSPF and associate it with the backbone area ID 0
`(config-router)#network 192.168.24.0/24 area 0`	Define the interface eth5 connected to PE2 to run OSPF and associate it with the backbone area ID 0
`(config-router)#exit`	Exit the router OSPF mode
`(config)#router ldp`	Enter Router mode for LDP
`(config-router)#exit`	Exit the Router mode for LDP
`(config)#interface eth1`	Enter interface mode for interface eth1
`(config-if)#label-switching`	Enable label-switching for interface eth1
`(config-if)#enable-ldp ipv4`	Enable LDP for IPv4 on interface eth1
`(config-if)#exit`	Exit the interface mode for eth1
`(config)#interface eth3`	Enter interface mode for interface eth3
`(config-if)#label-switching`	Enable label-switching for interface eth3
`(config-if)#enable-ldp ipv4`	Enable LDP for IPv4 on interface eth3
`(config-if)#exit`	Exit the interface mode for eth3
`(config)#interface eth4`	Enter interface mode for interface eth4
`(config-if)#label-switching`	Enable label-switching for interface eth4
`(config-if)#enable-ldp ipv4`	Enable LDP for IPv4 on interface eth4
`(config-if)#exit`	Exit the interface mode for eth4
`(config)#interface eth5`	Enter interface mode for interface eth5
`(config-if)#label-switching`	Enable label-switching for interface eth5
`(config-if)#enable-ldp ipv4`	Enable LDP for IPv4 on interface eth5
`(config-if)#exit`	Return to the global mode
`(config)#commit`	Apply commit
`(config)#exit`	exit the global mode

CE2 Configuration

`#configure terminal`	Enter configure mode
`(config)#interface eth2`	Enter Interface mode for eth2.
`(config-if)#ip address 10.0.22.1/24`	Assign IP address.
`(config-if)#exit`	Exit interface mode.
`(config)#interface lo`	Enter Interface mode for lo.
`(config-if)#ip address 200.1.1.2/24`	Assign IP address.
`(config-if)#exit`	Exit interface mode.
`(config)#router bgp 65002`	Enter BGP router mode.

© 2026 IP Infusion Inc. Proprietary

(config-router)#neighbor 10.0.22.2 remote-as 100	Configure an eBGP neighbor on PE1
(config-router)#address-family ipv4 unicast	Configure the address family ipv4 under router BGP
(config-router-af)#neighbor 10.0.22.2 activate	Activate eBGP neighbor on PE1
(config-router-af)#network 200.1.1.0/24	Announce the network 200.1.1.0/24 into BGP
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#exit	Exit the router BGP mode and return to the configure mode
(config)#commit	Apply commit
(config)#exit	Exit the configure mode

Additional Paths Capability Send/Receive/Send-Receive and Additional Paths Selection All/Best 2/Best 3

Configure RR to Send All And Pe1 To Receive All Additional Paths at Address Family Vpnv4 Level

RR

#configure terminal	Enter configure mode
(config)#router bgp 100	Enter BGP router mode
(config-router)#address-family vpnv4 unicast	Enter address family VPNv4 unicast mode
(config-router-af)#bgp additional-paths send	Configure RR to send additional paths to all the VPNv4 neighbors
(config-router-af)#bgp additional-paths select all	Configure RR to select all available paths to send to all VPNv4 neighbors
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#exit	Exit Router BGP and Configure mode.
(config)#commit	Apply commit
(config)#exit	Exit Router BGP and Configure mode.

PE1

#configure terminal	Enter configure mode
(config)#router bgp 100	Enter BGP router mode
(config-router)#address-family vpnv4 unicast	Enter address family VPNv4 unicast mode
(config-router-af)#bgp additional-paths receive	Configure PE1 to receive additional paths from all the VPNv4 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 return to the configure mode

© 2026 IP Infusion Inc. Proprietary

(config)#commit	Apply commit
(config)#exit	Exit the configure mode

Configure RR To Send All and Pe1 To Receive All Additional Paths at Neighbor Level Address Family VPNv4

RR

#configure terminal	Enter configure mode
(config)#router bgp 100	Enter BGP router mode
(config-router)#address-family vpnv4 unicast	Enter address family VPNv4 unicast mode
(config-router-af)#neighbor 192.168.11.2 additional-paths send	Configure RR to send additional paths to the VPNv4 neighbor on PE1
(config-router-af)#neighbor 192.168.11.2 advertise additional-paths all	Configure RR to advertise all available paths to the VPNv4 neighbor on PE1
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#exit	Exit Router BGP and Configure mode.
(config)#commit	Apply commit
(config)#exit	Exit Router BGP and Configure mode.

PE1

#configure terminal	Enter configure mode
(config)#router bgp 100	Enter BGP router mode
(config-router)#address-family vpnv4 unicast	Enter address family VPNv4 unicast mode
(config-router-af)#neighbor 192.168.11.1 additional-paths receive	Configure PE1 to receive additional paths from RR
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#exit	Exit Router BGP and Configure mode.
(config)#commit	Apply commit
(config)#exit	Exit Router BGP and Configure mode.

Validation

RR

```
#show ip bgp neighbors 192.168.11.2
BGP neighbor is 192.168.11.2, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 192.168.11.2
  BGP state = Established, up for 00:03:21
  Last read 00:03:21, 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
    Address family VPNv4 Unicast: advertised and received
```

```
  Received 90 messages, 1 notifications, 0 in queue
  Sent 94 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
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: VPNv4 Unicast
  BGP table version 49, neighbor version 49
  Index 1, Offset 0, Mask 0x2
  AF-dependant capabilities:
    Add-Path Send Capability : advertised
    Add-Path Receive Capability : received
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  4 announced prefixes

 Connections established 3; dropped 2
Local host: 192.168.11.1, Local port: 179
Foreign host: 192.168.11.2, Foreign port: 53977
Nexthop: 192.168.11.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:03:26, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)

#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: 100:1
*>i  100.1.1.0/24     192.168.11.2          0        100          0       65001 i
*>i  200.1.1.0        192.168.21.2          0        100          0       65002 i
* i  200.1.1.0        192.168.22.2          0        100          0       65002 i
* i  200.1.1.0        192.168.23.2          0        100          0       65002 i
* i  200.1.1.0        192.168.24.2          0        100          0       65002 i
 Announced routes count = 0
 Accepted routes count = 5

#show ip bgp vpnv4 all 200.1.1.0
Route Distinguisher: 100:1
  65002, (Received from a RR-client)
    192.168.21.2 from 192.168.21.2 (192.168.24.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal, best
      Extended Community: RT:200:1
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       192.168.11.2
      Last update: Mon Jan  9 05:27:09 2017

  65002, (Received from a RR-client)
    192.168.22.2 from 192.168.22.2 (192.168.24.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal
      Extended Community: RT:200:1
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       192.168.11.2
      Last update: Mon Jan  9 05:27:09 2017

  65002, (Received from a RR-client)
```

```
    192.168.23.2 from 192.168.23.2 (192.168.24.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal
      Extended Community: RT:200:1
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       192.168.11.2
      Last update: Mon Jan  9 05:27:09 2017

  65002, (Received from a RR-client)
    192.168.24.2 from 192.168.24.2 (192.168.24.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal
      Extended Community: RT:200:1
      rx path_id: -1      tx path_id: 3
      Advertised to non peer-group peers:
       192.168.11.2
      Last update: Mon Jan  9 05:27:09 2017
```

PE1

```
#show ip bgp neighbors 192.168.11.1
BGP neighbor is 192.168.11.1, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 10.12.48.36
  BGP state = Established, up for 00:46:03
  Last read 00:46:03, 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
    Address family VPNv4 Unicast: advertised and received
  Received 179 messages, 1 notifications, 0 in queue
  Sent 176 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 83, neighbor version 82
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: VPNv4 Unicast
  BGP table version 9, neighbor version 9
  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
  0 announced prefixes

 Connections established 3; dropped 2
Local host: 192.168.11.2, Local port: 53977
Foreign host: 192.168.11.1, Foreign port: 179
Nexthop: 192.168.11.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:46:08, due to BGP Notification sent
Notification Error Message: (Cease/Other Configuration Change.)

#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: 100:1 (Default for VRF vrf1)
*>   100.1.1.0/24     10.0.11.1            0        100          0       65001 i
*>i  200.1.1.0        192.168.24.2         0        100          0       65002 i
```

```
 Announced routes count = 1
 Accepted routes count = 1
Route Distinguisher: 100:1
*>i  200.1.1.0        192.168.24.2         0        100          0       65002 i
* i  200.1.1.0        192.168.23.2         0        100          0       65002 i
* i  200.1.1.0        192.168.22.2         0        100          0       65002 i
* i  200.1.1.0        192.168.21.2         0        100          0       65002 i
 Announced routes count = 0
 Accepted routes count = 4

#show ip bgp vpnv4 all 200.1.1.0
Route Distinguisher: 100:1 (Default for VRF vrf1)
  65002
    192.168.24.2 (metric 2) from 192.168.11.1 (192.168.24.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal, best
      Extended Community: RT:200:1
      Originator: 192.168.24.2, Cluster list: 10.12.48.36
      rx path_id: 0      tx path_id: 0
      Not advertised to any peer
      Last update: Mon Jan  9 05:28:04 2017

Route Distinguisher: 100:1
  65002
    192.168.24.2 (metric 2) from 192.168.11.1 (192.168.24.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal, best
      Extended Community: RT:200:1
      Originator: 192.168.24.2, Cluster list: 10.12.48.36
      rx path_id: 3      tx path_id: 1
      Not advertised to any peer
      Last update: Mon Jan  9 05:28:04 2017

  65002
    192.168.23.2 (metric 2) from 192.168.11.1 (192.168.24.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal
      Extended Community: RT:200:1
      Originator: 192.168.24.2, Cluster list: 10.12.48.36
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Jan  9 05:28:04 2017

  65002
    192.168.22.2 (metric 2) from 192.168.11.1 (192.168.24.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal
      Extended Community: RT:200:1
      Originator: 192.168.24.2, Cluster list: 10.12.48.36
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Jan  9 05:28:04 2017

  65002
    192.168.21.2 (metric 2) from 192.168.11.1 (192.168.24.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal
      Extended Community: RT:200:1
      Originator: 192.168.24.2, Cluster list: 10.12.48.36
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Jan  9 05:27:30 2017
```

Configure RR to Send Best 2 and Pe1 To Receive All Additional Paths at Address Family VPNv4 Level

RR

`#configure terminal`	Enter configure mode
`(config)#router bgp 100`	Enter BGP router mode

(config-router)#address-family vpnv4 unicast	Enter address family VPNv4 unicast mode
(config-router-af)#bgp additional-paths send	Configure RR to send additional paths to all the VPNv4 neighbors
(config-router-af)#bgp additional-paths select best 2	Configure RR to select best 2 available paths to send to all VPNv4 neighbors
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#exit	Exit Router BGP and Configure mode.
(config)#commit	Apply commit
(config)#exit	Exit Router BGP and Configure mode.

PE1

#configure terminal	Enter configure mode
(config)#router bgp 100	Enter BGP router mode
(config-router)#address-family vpnv4 unicast	Enter address family VPNv4 unicast mode
(config-router-af)#bgp additional-paths receive	Configure PE1 to receive additional paths from all the VPNv4 neighbors
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#exit	Exit Router BGP and Configure mode.
(config)#commit	Apply commit
(config)#exit	Exit Router BGP and Configure mode.

Configure RR to Send Best 2 and Pe1 To Receive All Additional Paths at Neighbor Level Address Family VPNv4

RR

#configure terminal	Enter configure mode
(config)#router bgp 100	Enter BGP router mode
(config-router)#address-family vpnv4 unicast	Enter address family VPNv4 unicast mode
(config-router-af)#neighbor 192.168.11.2 additional-paths send	Configure RR to send additional paths to the VPNv4 neighbor on PE1
(config-router-af)#neighbor 192.168.11.2 advertise additional-paths best 2	Configure RR to advertise best 2 out of all available paths to the VPNv4 neighbor on PE1
(config-router-af)#end	Return to the global mode
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#exit	Exit Router BGP and Configure mode.
(config)#commit	Apply commit
(config)#exit	Exit Router BGP and Configure mode.

© 2026 IP Infusion Inc. Proprietary

PE1

`#configure terminal`	Enter configure mode
`(config)#router bgp 100`	Enter BGP router mode
`(config-router)#address-family vpnv4 unicast`	Enter address family VPNv4 unicast mode
`(config-router-af)#neighbor 192.168.11.1 additional-paths receive`	Configure PE1 to receive additional paths from RR
`(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`(config-router)#exit`	Exit Router BGP and Configure mode.
`(config)#commit`	Apply commit
`(config)#exit`	Exit Router BGP and Configure mode.

Validation

RR

```
#show ip bgp vpnv4 all 200.1.1.0
Route Distinguisher: 100:1
  65002, (Received from a RR-client)
    192.168.21.2 from 192.168.21.2 (192.168.21.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal, best
      Extended Community: RT:200:1
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       192.168.11.2
      Last update: Wed Jan 11 06:03:50 2017

  65002, (Received from a RR-client)
    192.168.22.2 from 192.168.22.2 (192.168.21.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal
      Extended Community: RT:200:1
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       192.168.11.2
      Last update: Wed Jan 11 06:02:10 2017

  65002, (Received from a RR-client)
    192.168.23.2 from 192.168.23.2 (192.168.21.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal
      Extended Community: RT:200:1
      rx path_id: -1      tx path_id: 1
      Not advertised to any peer
      Last update: Wed Jan 11 06:02:10 2017

  65002, (Received from a RR-client)
    192.168.24.2 from 192.168.24.2 (192.168.21.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal
      Extended Community: RT:200:1
      rx path_id: -1      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:02:10 2017
```

PE1

```
#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: 100:1 (Default for VRF vrf1)
*>   100.1.1.0/24     10.0.11.1            0        100          0       65001 i
*>i  200.1.1.0        192.168.21.2         0        100          0       65002 i
 Announced routes count = 1
 Accepted routes count = 1
Route Distinguisher: 100:1
*>i  200.1.1.0        192.168.21.2         0        100          0       65002 i
* i  200.1.1.0        192.168.22.2         0        100          0       65002 i
 Announced routes count = 0
 Accepted routes count = 2

#show ip bgp vpnv4 all 200.1.1.0
Route Distinguisher: 100:1 (Default for VRF vrf1)
  65002
    192.168.21.2 (metric 2) from 192.168.11.1 (192.168.21.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal, best
      Extended Community: RT:200:1
      Originator: 192.168.21.2, Cluster list: 10.12.48.36
      rx path_id: 0      tx path_id: 0
      Not advertised to any peer
      Last update: Wed Jan 11 06:03:43 2017

Route Distinguisher: 100:1
  65002
    192.168.21.2 (metric 2) from 192.168.11.1 (192.168.21.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal, best
      Extended Community: RT:200:1
      Originator: 192.168.21.2, Cluster list: 10.12.48.36
      rx path_id: 2      tx path_id: 1
      Not advertised to any peer
      Last update: Wed Jan 11 06:03:43 2017

  65002
    192.168.22.2 (metric 2) from 192.168.11.1 (192.168.21.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal
      Extended Community: RT:200:1
      Originator: 192.168.21.2, Cluster list: 10.12.48.36
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:02:05 2017
```

Configure RR to Send Best 3 and Pe1 to Receive All Additional Paths at Address Family VPNv4 Level

RR

`#configure terminal`	Enter configure mode
`(config)#router bgp 100`	Enter BGP router mode
`(config-router)#address-family vpnv4 unicast`	Enter address family VPNv4 unicast mode
`(config-router-af)#bgp additional-paths send`	Configure RR to send additional paths to all the VPNv4 neighbors
`(config-router-af)#bgp additional-paths select best 3`	Configure RR to select best 3 available paths to send to all VPNv4 neighbors
`(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`(config-router)#exit`	Exit Router BGP and Configure mode.

(config)#commit	Apply commit
(config)#exit	Exit Router BGP and Configure mode.

PE1

#configure terminal	Enter configure mode
(config)#router bgp 100	Enter BGP router mode
(config-router)#address-family vpnv4 unicast	Enter address family VPNv4 unicast mode
(config-router-af)#bgp additional-paths receive	Configure PE1 to receive additional paths from all the VPNv4 neighbors
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#exit	Exit Router BGP and Configure mode.
(config)#commit	Apply commit
(config)#exit	Exit Router BGP and Configure mode.

Configure RR to Send Best 3 and Pe1 to Receive All Additional Paths at Neighbor Level Address Family VPNv4

RR

#configure terminal	Enter configure mode
(config)#router bgp 100	Enter BGP router mode
(config-router)#address-family vpnv4 unicast	Enter address family VPNv4 unicast mode
(config-router-af)#neighbor 192.168.11.2 additional-paths send	Configure RR to send additional paths to the VPNv4 neighbor on PE1
(config-router-af)#neighbor 192.168.11.2 advertise additional-paths best 3	Configure RR to advertise best 3 out of all available paths to the VPNv4 neighbor on PE1
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.
(config-router)#exit	Exit Router BGP and Configure mode.
(config)#commit	Apply commit
(config)#exit	Exit Router BGP and Configure mode.

PE1

#configure terminal	Enter configure mode
(config)#router bgp 100	Enter BGP router mode
(config-router)#address-family vpnv4 unicast	Enter address family VPNv4 unicast mode
(config-router-af)#neighbor 192.168.11.1 additional-paths receive	Configure PE1 to receive additional paths from RR
(config-router-af)#exit-address-family	Exit Address Family mode and return to Router mode.

© 2026 IP Infusion Inc. Proprietary

`(config-router)#exit`	Exit Router BGP and Configure mode.
`(config)#commit`	Apply commit
`(config)#exit`	Exit Router BGP and Configure mode.

Validation

RR

```
#show ip bgp vpnv4 all 200.1.1.0
Route Distinguisher: 100:1
  65002, (Received from a RR-client)
    192.168.21.2 from 192.168.21.2 (192.168.21.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal, best
      Extended Community: RT:200:1
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       192.168.11.2
      Last update: Wed Jan 11 06:03:50 2017

  65002, (Received from a RR-client)
    192.168.22.2 from 192.168.22.2 (192.168.21.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal
      Extended Community: RT:200:1
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       192.168.11.2
      Last update: Wed Jan 11 06:02:10 2017

  65002, (Received from a RR-client)
    192.168.23.2 from 192.168.23.2 (192.168.21.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal
      Extended Community: RT:200:1
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       192.168.11.2
      Last update: Wed Jan 11 06:02:10 2017

  65002, (Received from a RR-client)
    192.168.24.2 from 192.168.24.2 (192.168.21.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal
      Extended Community: RT:200:1
      rx path_id: -1      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:02:10 2017
```

PE1

```
#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: 100:1 (Default for VRF vrf1)
*>   100.1.1.0/24     10.0.11.1            0        100          0       65001 i
*>i  200.1.1.0        192.168.23.2         0        100          0       65002 i
 Announced routes count = 1
 Accepted routes count = 1
Route Distinguisher: 100:1
*>i  200.1.1.0        192.168.23.2         0        100          0       65002 i
* i  200.1.1.0        192.168.22.2         0        100          0       65002 i
* i  200.1.1.0        192.168.21.2         0        100          0       65002 i
 Announced routes count = 0
```

```
 Accepted routes count = 3

#show ip bgp vpnv4 all 200.1.1.0
Route Distinguisher: 100:1 (Default for VRF vrf1)
  65002
    192.168.23.2 (metric 2) from 192.168.11.1 (192.168.21.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal, best
      Extended Community: RT:200:1
      Originator: 192.168.21.2, Cluster list: 10.12.48.36
      rx path_id: 0      tx path_id: 0
      Not advertised to any peer
      Last update: Wed Jan 11 06:24:47 2017

Route Distinguisher: 100:1
  65002
    192.168.23.2 (metric 2) from 192.168.11.1 (192.168.21.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal, best
      Extended Community: RT:200:1
      Originator: 192.168.21.2, Cluster list: 10.12.48.36
      rx path_id: 1      tx path_id: 1
      Not advertised to any peer
      Last update: Wed Jan 11 06:24:47 2017

  65002
    192.168.22.2 (metric 2) from 192.168.11.1 (192.168.21.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal
      Extended Community: RT:200:1
      Originator: 192.168.21.2, Cluster list: 10.12.48.36
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:24:47 2017

  65002
    192.168.21.2 (metric 2) from 192.168.11.1 (192.168.21.2)
      Origin IGP, metric 0, localpref 100, label    24320, valid, internal
      Extended Community: RT:200:1
      Originator: 192.168.21.2, Cluster list: 10.12.48.36
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:24:00 2017
```

Additional Paths for 6PE

This section contains the configuration of additional paths group best for 6PE.

Topology

Figure 84. Additional paths group best for 6PE

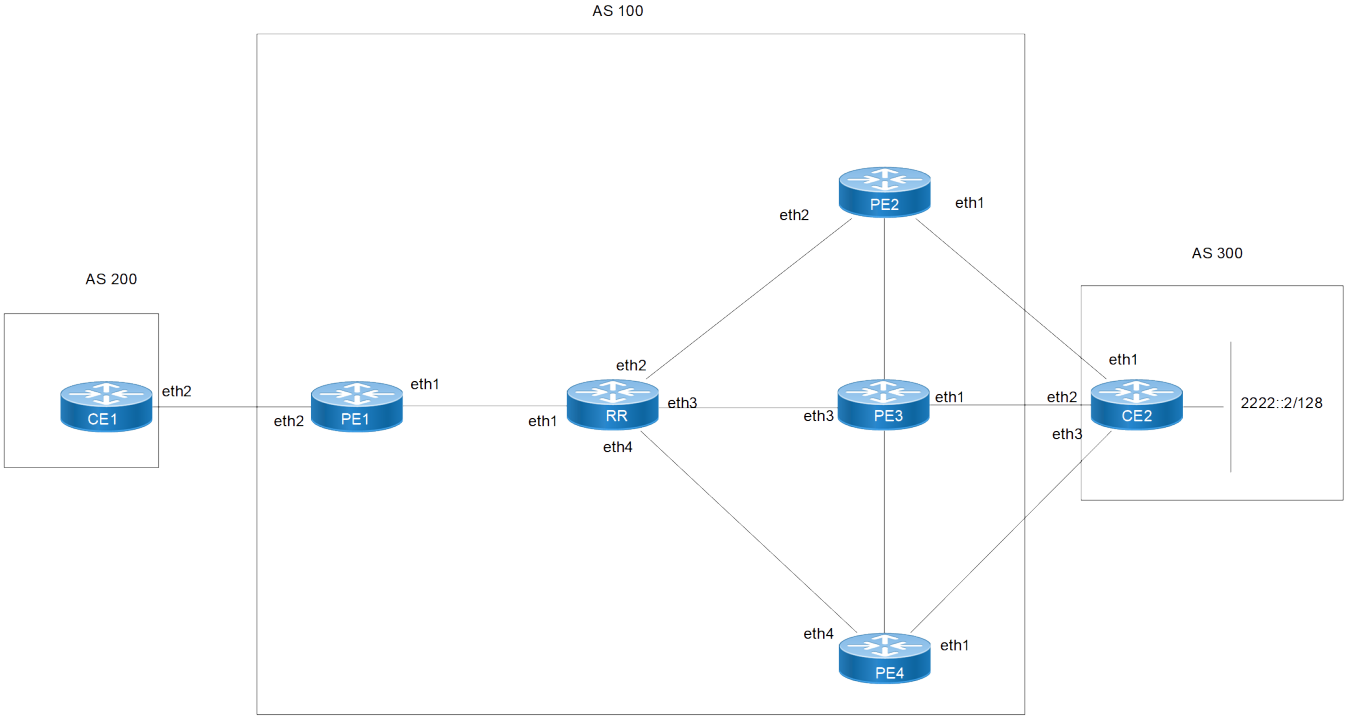

Configuration

CE1: Interface

`CE1#configure terminal`	Enter configure mode.
`CE1(config)#hostname CE1`	Configure the hostname to CE1
`CE1(config)#interface eth2`	Enter the Interface mode for interface eth2.
`CE1(config-if)#ipv6 address 1001::1/64`	Configure IPV6 address for interface eth2
`(config)#commit`	Apply commit
`(config)#exit`	Exit Router BGP and Configure mode.

CE1: BGP

`CE1(config)#router bgp 200`	Enter BGP router mode
`CE1(config-router)#neighbor 1001::2 remote-as 100`	Configure BGP ipv6 neighbor on PE1
`CE1(config-router)#address-family ipv6 unicast`	Enter the address family ipv6
`CE1(config-router-af)#neighbor 1001::2 activate`	Activate neighbor under address family ipv6
`CE1(config-router-af)#exit-address-family`	Exit address family ipv6 unicast mode
`CE1(config-router)#exit`	Exit BGP router mode
`CE1(config)#commit`	Apply commit
`CE1#clear bgp 1001::2`	Clear BGP peer PE1

© 2026 IP Infusion Inc. Proprietary

PE1: Interface

`PE1#configure terminal`	Enter configure mode.
`PE1(config)#hostname PE1`	Configure the hostname to PE1
`PE1(config)#interface eth2`	Enter the Interface mode for interface eth2
`PE1(config-if)#ipv6 address 1001::2/64`	Configure the IPv6 address for interface eth2
`PE1(config-if)#exit`	Exit interface mode
`PE1(config)#interface lo`	Enter the Interface mode for the loopback interface.
`PE1(config-if)#ip address 1.1.1.1/32`	Configure IPV4 address for loopback interface.
`PE1(config-if)#exit`	Exit interface mode
`PE1(config)#interface eth1`	Enter the Interface mode for interface eth1
`PE1(config-if)#ip address 10.1.1.1/24`	Configure the IP address for interface eth1
`(config-if)#exit`	Exit the interface level
`(config)#commit`	Apply commit
`(config)#exit`	Exit.

PE1: OSPF

`PE1(config)#router ospf 1`	Configure an OSPF instance 1
`PE1(config-router)#network 10.1.1.0/24 area 0`	Advertise the network towards P in Area 0
`PE1(config-router)#network 1.1.1.1/32 area 0`	Advertise the Loopback network in Area 0
`PE1(config-router)#exit`	Exit OSPF router mode
`PE1(config)#commit`	Apply commit
`PE1(config)#exit`	Exit the global mode

PE1: MPLS

`PE1(config)#router ldp`	Enable LDP Process
`PE1(config-router)#transport-address ipv4 1.1.1.1`	Configure Transport address for LDP
`PE1(config-router)#exit`	Exit router ldp mode
`PE1(config)#interface eth1`	Enter the interface mode for interface eth1
`PE1(config-if)#label-switching`	Enable label-switching on interface eth1
`PE1(config-if)#enable-ldp ipv4`	Enable LDP for IPv4 on interface eth1
`PE1(config-if)#exit`	Exit the interface mode
`PE1(config)#commit`	Apply commit
`PE1(config)#exit`	Exit the global mode

PE1: BGP

`PE1(config)#router bgp 100`	Configure router bgp mode

© 2026 IP Infusion Inc. Proprietary

`PE1(config)#bgp router-id 1.1.1.1`	Configure BGP router-id
`PE1(config-router)#neighbor 2.2.2.2 remote-as 100`	Configure BGP IPv4 neighbor on PE2
`PE1(config-router)#neighbor 2.2.2.2 update-source 1.1.1.1`	Configure the routing update source for RR as PE1's Loopback
`PE1(config-router)#address-family ipv6 labeled-unicast`	Enter BGP address-family ipv6 labeled unicast mode
`PE1(config-router-af)#neighbor 2.2.2.2 activate`	Activate the neighbor RR for address-family ipv6 labeled-unicast unicast
`PE1(config-router-af)#exit-address-family`	Exit address-family ipv6 labeled-unicast mode
`PE1(config-router)#address-family ipv6 unicast`	Enter BGP address-family ipv6 unocast mode
`PE1(config-router)#neighbor 1001::1 remote-as 200`	Configure BGP IPv6 neighbor on CE1
`PE1(config-router-af)#exit-address-family`	Exit address-family ipv6 unicast mode
`PE1(config-router)#exit`	Exit router bgp mode
`PE1(config)#exit`	Exit configure mode
`PE1(config)#commit`	Apply commit

RR: Interface

`RR#configure terminal`	Enter configure mode.
`RR(config)#hostname RR`	Configure hostname as RR
`RR(config)#interface eth1`	Enter interface mode for interface eth1
`RR(config-if)#ip address 10.1.1.2/24`	Configure IPv4 address for interface eth1
`RR(config-if)#exit`	Exit interface mode
`RR(config)#interface eth2`	Enter interface mode for interface eth2
`RR(config-if)#ip address 20.1.1.2/24`	Configure IPv4 address for interface eth2
`RR(config-if)#exit`	Exit interface mode
`RR(config)#interface eth3`	Enter interface mode for interface eth3
`RR(config-if)#ip address 30.1.1.2/24`	Configure IPv4 address for interface eth3
`RR(config-if)#exit`	Exit interface mode
`RR(config)#interface eth4`	Enter interface mode for interface eth4
`RR(config-if)#ip address 40.1.1.2/24`	Configure IPv4 address for interface eth4
`RR(config-if)#exit`	Exit interface mode
`RR(config)#interface lo`	Enter Interface mode for the loopback interface.
`RR(config-if)#ip address 2.2.2.2/32`	Configure IPV4 address for loopback interface.
`RR(config-if)#exit`	Exit Interface mode.
`PE1(config)#commit`	Apply commit

© 2026 IP Infusion Inc. Proprietary

RR: OSPF

`RR(config)#router ospf 1`	Configure router OSPF instance 1
`RR(config-router)#network 10.1.1.0/24 area 0`	Advertise the network towards PE1 in area 0
`RR(config-router)#network 20.1.1.0/24 area 0`	Advertise the network towards PE2 in area 0
`RR(config-router)#network 30.1.1.0/24 area 0`	Advertise the network towards PE3 in area 0
`RR(config-router)#network 40.1.1.0/24 area 0`	Advertise the network towards PE4 in area 0
`RR(config-router)#network 2.2.2.2/32 area 0`	Advertise the Loopback network in area 0
`RR(config-router)#exit`	Exit the router ospf mode
`RR(config)#commit`	Apply commit

RR: MPLS

`RR(config)#router ldp`	Enable LDP Process
`RR(config-router)#transport-address ipv4 2.2.2.2`	Configure Transport address for LDP
`RR(config-router)#exit`	Exit router ldp mode
`RR(config)#interface eth1`	Enter the interface mode for interface eth1
`RR(config-if)#label-switching`	Enable label-switching on interface eth1
`RR(config-if)#enable-ldp ipv4`	Enable LDP for IPv4 on interface eth1
`RR(config-if)#exit`	Exit interface mode
`RR(config)#interface eth2`	Enter the interface mode for interface eth2
`RR(config-if)#label-switching`	Enable label-switching on interface eth2
`RR(config-if)#enable-ldp ipv4`	Enable LDP for IPv4 on interface eth2
`RR(config-if)#exit`	Exit interface mode
`RR(config)#interface eth3`	Enter the interface mode for interface eth3
`RR(config-if)#label-switching`	Enable label-switching on interface eth3
`RR(config-if)#enable-ldp ipv4`	Enable LDP for IPv4 on interface eth3
`RR(config-if)#exit`	Exit interface mode
`RR(config)#interface eth4`	Enter the interface mode for interface eth4
`RR(config-if)#label-switching`	Enable label-switching on interface eth4
`RR(config-if)#enable-ldp ipv4`	Enable LDP for IPv4 on interface eth4
`RR(config-if)#exit`	Exit interface mode
`RR(config)#commit`	Apply commit
`RR(config)#exit`	Exit the configure terminal mode

RR: BGP

`RR(config)#router bgp 100`	Configure router bgp mode
`RR(config)#bgp router-id 2.2.2.2`	Configure BGP router-id

`RR(config-router)#neighbor 1.1.1.1 remote-as 100`	Configure BGP neighbor on PE1
`RR(config-router)#neighbor 1.1.1.1 update-source 2.2.2.2`	Configure the routing update source for PE1 as RR's Loopback
`RR(config-router)#neighbor 3.3.3.3 remote-as 100`	Configure BGP neighbor on PE2
`RR(config-router)#neighbor 3.3.3.3 update-source 2.2.2.2`	Configure the routing update source for PE2 as RR's Loopback
`RR(config-router)#neighbor 4.4.4.4 remote-as 100`	Configure BGP neighbor on PE3
`RR(config-router)#neighbor 4.4.4.4 update-source 2.2.2.2`	Configure the routing update source for PE3 as RR's Loopback
`RR(config-router)#neighbor 5.5.5.5 remote-as 100`	Configure BGP neighbor on PE4
`RR(config-router)#neighbor 5.5.5.5 update-source 2.2.2.2`	Configure the routing update source for PE4 as RR's Loopback
`RR(config-router)#address-family ipv6 labeled-unicast`	Enter BGP address-family ipv6 labeled unicast mode
`RR(config-router-af)#neighbor 1.1.1.1 activate`	Activate the neighbor PE1 for address-family ipv6 lu unicast
`RR(config-router-af)#neighbor 1.1.1.1 route-reflector-client`	Configure PE1 as route-reflector-client
`RR(config-router-af)#neighbor 3.3.3.3 activate`	Activate the neighbor PE2 for address-family ipv6 lu unicast
`RR(config-router-af)#neighbor 3.3.3.3 route-reflector-client`	Configure PE2 as route-reflector-client
`RR(config-router-af)#neighbor 4.4.4.4 activate`	Activate the neighbor PE3 for address-family ipv6 lu unicast
`RR(config-router-af)#neighbor 4.4.4.4 route-reflector-client`	Configure PE3 as route-reflector-client
`RR(config-router-af)#neighbor 5.5.5.5 activate`	Activate the neighbor PE4 for address-family ipv6 lu unicast
`RR(config-router-af)#neighbor 5.5.5.5 route-reflector-client`	Configure PE4 as route-reflector-client
`RR(config-router-af)#exit-address-family`	Exit BGP address-family ipv6 lu unicast mode
`RR(config-router)#exit`	Exit router bgp mode
`RR(config)#commit`	Apply commit
`RR(config)#exit`	Exit configure mode

PE2: Interface

`PE2#configure terminal`	Enter configure mode.
`PE2(config)#hostname PE2`	Configure the hostname to PE2
`PE2(config)#interface eth1`	Enter the Interface mode for interface eth1
`PE2(config-if)#ipv6 address 2001::1/64`	Configure the IPv6 address for interface eth1
`PE2(config-if)#exit`	Exit interface mode

© 2026 IP Infusion Inc. Proprietary

PE2(config)#interface lo	Enter the Interface mode for the loopback interface.
PE2(config-if)#ip address 3.3.3.3/32	Configure IPV4 address for loopback interface.
PE2(config-if)#exit	Exit interface mode
PE2(config)#interface eth2	Enter the Interface mode for interface eth2
PE2(config-if)#ip address 20.1.1.1/24	Configure the IP address for interface eth2
PE2(config-if)#exit	Exit configure mode
PE2(config)#commit	Apply commit

PE2: OSPF

PE2(config)#router ospf 1	Configure an OSPF instance 1
PE2(config-router)#network 20.1.1.0/24 area 0	Advertise the network towards RR in Area 0
PE2(config-router)#network 3.3.3.3/32 area 0	Advertise the Loopback network in Area 0
PE2(config-router)#exit	Exit OSPF router mode
PE2(config)#commit	Apply commit

PE2: MPLS

PE2(config)#router ldp	Enable LDP Process
PE2(config-router)#transport-address ipv4 3.3.3.3	Configure Transport address for LDP
PE2(config-router)#exit	Exit router ldp mode
PE2(config)#interface eth2	Enter the interface mode for interface eth2
PE2(config-if)#label-switching	Enable label-switching on interface eth2
PE2(config-if)#enable-ldp ipv4	Enable LDP for IPv4 on interface eth2
PE2(config-if)#exit	Exit interface mode
PE2(config)#commit	Apply commit

PE2: BGP

PE2(config)#router bgp 100	Configure router bgp mode
PE2(config)#bgp router-id 3.3.3.3	Configure BGP router-id
PE2(config-router)#neighbor 2.2.2.2 remote-as 100	Configure BGP IPv4 neighbor on RR
PE2(config-router)#neighbor 2.2.2.2 update-source 3.3.3.3	Configure the routing update source for RR as PE2's Loopback
PE2(config-router)#address-family ipv6 labeled-unicast	Enter BGP address-family ipv6 labeled-unicast mode
PE2(config-router-af)#neighbor 2.2.2.2 activate	Activate the neighbor RR for address-family ipv6 labeled-unicast
PE2(config-router-af)#exit-address-family	Exit address-family ipv6 labeled-unicast mode
PE2(config-router)#address-family ipv6 unicast	Enter BGP address-family ipv6 unicast mode
PE2(config-router)#neighbor 2001::2 remote-as 300	Configure BGP IPv6 neighbor on CE2

© 2026 IP Infusion Inc. Proprietary

PE2(config-router-af)#exit-address-family	Exit address-family ipv6 unicast mode
PE2(config-router)#exit	Exit router bgp mode
PE2(config)#commit	Apply commit
PE2(config)#exit	Exit configure mode
PE2#clear bgp ipv6 unicast 2001::2 unicast	Clear BGP neighbor CE2

PE3: Interface

PE3#configure terminal	Enter configure mode.
PE3(config)#hostname PE3	Configure the hostname to PE3
PE3(config)#interface eth1	Enter the Interface mode for interface eth1
PE3(config-if)#ipv6 address 3001::1/64	Configure the IPv6 address for interface eth1
PE3(config-if)#exit	Exit interface mode
PE3(config)#interface lo	Enter the Interface mode for the loopback interface.
PE3(config-if)#ip address 4.4.4.4/32	Configure IPV4 address for loopback interface.
PE3(config-if)#exit	Exit interface mode
PE3(config)#interface eth3	Enter the Interface mode for interface eth3
PE3(config-if)#ip address 30.1.1.1/24	Configure the IP address for interface eth3
PE3(config-if)#exit	Exit interface mode
PE3(config)#commit	Apply commit

PE3: OSPF

PE3(config)#router ospf 1	Configure an OSPF instance 1
PE3(config-router)#network 30.1.1.0/24 area 0	Advertise the network towards RR in Area 0
PE3(config-router)#network 4.4.4.4/32 area 0	Advertise the Loopback network in Area 0
PE3(config-router)#exit	Exit OSPF router mode
PE3(config)#commit	Apply commit

PE3: MPLS

PE3(config)#router ldp	Enable LDP Process
PE3(config-router)#transport-address ipv4 4.4.4.4	Configure Transport address for LDP
PE3(config-router)#exit	Exit router ldp mode
PE3(config)#interface eth3	Enter the interface mode for interface eth3
PE3(config-if)#label-switching	Enable label-switching on interface eth3
PE3(config-if)#enable-ldp ipv4	Enable LDP for IPv4 on interface eth3
PE3(config-if)#exit	Exit interface mode
PE3(config)#commit	Apply commit

© 2026 IP Infusion Inc. Proprietary

PE3: BGP

`PE3(config)#router bgp 100`	Configure router bgp mode
`PE3(config)#bgp router-id 4.4.4.4`	Configure BGP router-id
`PE3(config-router)#neighbor 2.2.2.2 remote-as 100`	Configure BGP IPv4 neighbor on RR
`PE3(config-router)#neighbor 2.2.2.2 update-source 4.4.4.4`	Configure the routing update source for RR as PE3's Loopback
`PE3(config-router)#address-family ipv6 labeled-unicast`	Enter BGP address-family ipv6 labeled-unicast mode
`PE3(config-router-af)#neighbor 2.2.2.2 activate`	Activate the neighbor RR for address-family ipv6 labeled-unicast
`PE3(config-router-af)#exit-address-family`	Exit address-family ipv6 labeled-unicast mode
`PE3(config-router)#address-family ipv6 unicast`	Enter BGP address-family ipv6 unicast mode
`PE3(config-router)#neighbor 3001::2 remote-as 300`	Configure BGP IPv6 neighbor on CE2
`PE3(config-router-af)#exit-address-family`	Exit address-family ipv6 unicast mode
`PE3(config-router)#exit`	Exit router bgp mode
`PE3(config)#commit`	Apply commit
`PE3(config)#exit`	Exit configure mode

PE4: Interface

`PE4#configure terminal`	Enter configure mode.
`PE4(config)#hostname PE4`	Configure the hostname to PE4
`PE4(config)#interface eth1`	Enter the Interface mode for interface eth1
`PE4(config-if)#ipv6 address 4001::1/64`	Configure the IPv6 address for interface eth1
`PE4(config-if)#exit`	Exit interface mode
`PE4(config)#interface lo`	Enter the Interface mode for the loopback interface.
`PE4(config-if)#ip address 5.5.5.5/32`	Configure IPV4 address for loopback interface.
`PE4(config-if)#exit`	Exit interface mode
`PE4(config)#interface eth4`	Enter the Interface mode for interface eth4
`PE4(config-if)#ip address 40.1.1.1/24`	Configure the IP address for interface eth4
`PE4(config-if)#exit`	Exit interface mode
`PE4(config)#commit`	Apply commit

PE4: OSPF

`PE4(config)#router ospf 1`	Configure an OSPF instance 1
`PE4(config-router)#network 40.1.1.0/24 area 0`	Advertise the network towards RR in Area 0
`PE4(config-router)#network 5.5.5.5/32 area 0`	Advertise the Loopback network in Area 0
`PE4(config-router)#exit`	Exit OSPF router mode
`PE4(config)#commit`	Apply commit

© 2026 IP Infusion Inc. Proprietary

PE4: MPLS

`PE4(config)#router ldp`	Enable LDP Process
`PE4(config-router)#transport-address ipv4 5.5.5.5`	Configure Transport address for LDP
`PE4(config-router)#exit`	Exit router ldp mode
`PE4(config)#interface eth4`	Enter the interface mode for interface eth4
`PE4(config-if)#label-switching`	Enable label-switching on interface eth4
`PE4(config-if)#enable-ldp ipv4`	Enable LDP for IPv4 on interface eth4
`PE4(config-if)#exit`	Exit interface mode
`PE4(config)#commit`	Apply commit

PE4: BGP

`PE4(config)#router bgp 100`	Configure router bgp mode
`PE4(config)#bgp router-id 5.5.5.5`	Configure BGP router-id
`PE4(config-router)#neighbor 2.2.2.2 remote-as 100`	Configure BGP IPv4 neighbor on RR
`PE4(config-router)#neighbor 2.2.2.2 update-source 5.5.5.5`	Configure the routing update source for RR as PE4's Loopback
`PE4(config-router)#address-family ipv6 labeled-unicast`	Enter BGP address-family ipv6 labeled-unicast mode
`PE4(config-router-af)#neighbor 2.2.2.2 activate`	Activate the neighbor RR for address-family ipv6 labeled-unicast
`PE4(config-router-af)#exit-address-family`	Exit address-family ipv6 labeled-unicast mode
`PE4(config-router)#address-family ipv6 unicast`	Enter BGP address-family ipv6 unicast mode
`PE4(config-router)#neighbor 4001::2 remote-as 300`	Configure BGP IPv6 neighbor on CE2
`PE4(config-router-af)#exit-address-family`	Exit address-family ipv6 unicast mode
`PE4(config-router)#exit`	Exit router bgp mode
`PE4(config)#commit`	Apply commit
`PE4(config)#exit`	Exit configure mode

CE2: Interface

`CE2#configure terminal`	Enter configure mode.
`CE2(config)#hostname CE2`	Configure the hostname to CE2
`CE2(config)#interface eth1`	Enter the Interface mode for interface eth1
`CE2(config-if)#ipv6 address 2001::2/64`	Configure IPV6 address for interface eth1
`CE2(config-if)#exit`	Exit Interface mode.
`CE2(config)#interface eth2`	Enter the Interface mode for interface eth2
`CE2(config-if)#ipv6 address 3001::2/64`	Configure IPV6 address for interface eth2
`CE2(config-if)#exit`	Exit Interface mode.

© 2026 IP Infusion Inc. Proprietary

CE2(config)#interface eth3	Enter the Interface mode for interface eth3
CE2(config-if)#ipv6 address 4001::2/64	Configure IPV6 address for interface eth3
CE2(config-if)#exit	Exit Interface mode.
CE2(config)#interface lo	Enter the Interface mode for interface Loopback
CE2(config-if)#ipv6 address 2222::2/128	Configure IPV6 address for interface Loopback
CE2(config-if)#exit	Exit Interface mode.
CE2(config)#commit	Apply commit

CE2: BGP

CE2(config)#router bgp 300	Enter BGP router mode
CE2(config-router)#neighbor 2001::1 remote-as 100	Configure BGP ipv6 neighbor on PE2
CE2(config-router)#neighbor 3001::1 remote-as 100	Configure BGP ipv6 neighbor on PE3
CE2(config-router)#neighbor 4001::1 remote-as 100	Configure BGP ipv6 neighbor on PE4
CE2(config-router)#address-family ipv6 unicast	Enter the address family ipv6
CE2(config-router-af)#neighbor 2001::1 activate	Activate neighbor under address family ipv6
CE2(config-router-af)#neighbor 3001::1 activate	Activate neighbor under address family ipv6
CE2(config-router-af)#neighbor 4001::1 activate	Activate neighbor under address family ipv6
CE2(config-router-af)#network 2222::2/128	Advertise the IPv6 prefix of Loopback into BGP
CE2(config-router-af)#exit-address-family	Exit address family ipv6 unicast mode
CE2(config-router)#exit	Exit BGP router mode
CE2(config)#commit	Apply commit
CE2(config)#exit	Exit configure mode

BGP Addpath Capability for ipv6 labeled-unicast Unicast Address-family

RR

RR(config)#router bgp 100	Configure router bgp mode
RR(config-router)#address-family ipv6 labeled-unicast	Enter BGP address-family ipv6 labeled-unicast mode
RR(config-router-af)#bgp additional-paths select all	Configure RR to select all addpaths
RR(config-router-af)#bgp additional-paths send	Configure BGP addpath send capability under address-family ipv6 labeled-unicast
(config-router-af)#exit-address-family	Exit bgp address-family mode.
(config-router)#exit	Exit Router BGP mode.
(config)#commit	Apply commit
(config)#exit	Exit Configure mode.

© 2026 IP Infusion Inc. Proprietary

PE1

`PE1(config)#router bgp 100`	Configure router bgp mode
`PE1(config-router)#address-family ipv6 labeled-unicast`	Enter BGP address-family ipv6 labeled-unicast mode
`PE1(config-router-af)#neighbor 2.2.2.2 additional-paths receive`	Configure BGP addpath receive capability for the 6pe-neighbor RR under address-family ipv6 labeled-unicast
`PE1(config-router-af)#neighbor 2.2.2.2 additional-paths receive`	Configure BGP addpath receive capability for the 6pe-neighbor RR under address-family ipv6 labeled-unicast
`PE1(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`PE1(config-router)#exit`	Exit Router BGP mode.
`PE1(config)#commit`	Apply commit
`PE1(config)#exit`	Exit Configure mode.

PE2

`PE2(config)#router bgp 100`	Configure router bgp mode
`PE2(config-router)#address-family ipv6 labeled-unicast`	Enter BGP address-family ipv6 labeled-unicast mode
`PE2(config-router-af)#bgp additional-paths receive`	Configure BGP addpath receive capability under address-family ipv6 labeled-unicast
`(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`(config-router)#exit`	Exit Router BGP mode.
`(config)#commit`	Apply commit
`(config)#exit`	Exit Configure mode.

PE3

`PE3(config)#router bgp 100`	Configure router bgp mode
`PE3(config-router)#address-family ipv6 labeled-unicast`	Enter BGP address-family ipv6 labeled-unicast mode
`PE3(config-router-af)#bgp additional-paths receive`	Configure BGP addpath receive capability under address-family ipv6 labeled-unicast
`PE3(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`PE3(config-router)#exit`	Exit Router BGP mode.
`PE3(config)#commit`	Apply commit
`PE3(config)#exit`	Exit Configure mode.

© 2026 IP Infusion Inc. Proprietary

PE4

`PE4(config)#router bgp 100`	Configure router bgp mode
`PE4(config-router)#address-family ipv6 labeled-unicast`	Enter BGP address-family ipv6 labeled-unicast mode
`PE4(config-router-af)#bgp additional-paths receive`	Configure BGP addpath receive capability under address-family ipv6 labeled-unicast
`PE4(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`PE4(config-router)#exit`	Exit Router BGP mode.
`PE4(config)#commit`	Apply commit
`PE4(config)#exit`	Exit Configure mode.

Validation

RR

Add-Path Send Capability Advertised to and Received From 6pe Peers:

```
RR#show bgp neighbors 1.1.1.1
BGP neighbor is 1.1.1.1, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 1.1.1.1
  BGP state = Established, up for 00:03:03
  Last read 00:03:03, 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
    Address family IPv6 Labeled-Unicast Unicast: advertised and received
  Received 35 messages, 0 notifications, 0 in queue
  Sent 41 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is 2.2.2.2
 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)
  0 accepted prefixes
  0 announced prefixes

 For address family: IPv6 Labeled-Unicast Unicast
  BGP table version 6, neighbor version 6
  Index 1, Offset 0, Mask 0x2
  AF-dependant capabilities:
    Add-Path Send Capability : advertised
    Add-Path Receive Capability : received
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  3 announced prefixes

 Connections established 2; dropped 1
Local host: 2.2.2.2, Local port: 57154
Foreign host: 1.1.1.1, Foreign port: 179
Nexthop: 2.2.2.2
Nexthop global: ::
Nexthop local: ::
```

```
BGP connection: non shared network
Last Reset: 00:03:08, due to BGP Notification sent
Notification Error Message: (Cease/Other Configuration Change.)
RR#


RR#show bgp neighbors 3.3.3.3
BGP neighbor is 3.3.3.3, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 3.3.3.3
  BGP state = Established, up for 00:05:20
  Last read 00:05:20, 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 Labeled-Unicast Unicast: advertised and received
  Received 41 messages, 0 notifications, 0 in queue
  Sent 44 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is 2.2.2.2
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: IPv6 Labeled-Unicast Unicast
  BGP table version 6, neighbor version 6
  Index 2, Offset 0, Mask 0x4
  AF-dependant capabilities:
    Add-Path Send Capability : advertised
    Add-Path Receive Capability : received
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  2 announced prefixes

 Connections established 2; dropped 1
Local host: 2.2.2.2, Local port: 179
Foreign host: 3.3.3.3, Foreign port: 32841
Nexthop: 2.2.2.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:05:25, due to BGP Notification sent
Notification Error Message: (Cease/Other Configuration Change.)
RR#


RR#show bgp neighbors 4.4.4.4
BGP neighbor is 4.4.4.4, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 4.4.4.4
  BGP state = Established, up for 00:05:47
  Last read 00:05:47, 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 Labeled-Unicast Unicast: advertised and received
  Received 42 messages, 0 notifications, 0 in queue
  Sent 46 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is 2.2.2.2
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 3, Offset 0, Mask 0x8
```

```
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: IPv6 Labeled-Unicast Unicast
  BGP table version 6, neighbor version 6
  Index 3, Offset 0, Mask 0x8
  AF-dependant capabilities:
    Add-Path Send Capability : advertised
    Add-Path Receive Capability : received
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  2 announced prefixes

 Connections established 2; dropped 1
Local host: 2.2.2.2, Local port: 179
Foreign host: 4.4.4.4, Foreign port: 42015
Nexthop: 2.2.2.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:05:47, due to BGP Notification received
Notification Error Message: (Cease/Administratively Reset.)
RR#


RR#show bgp neighbors 5.5.5.5
BGP neighbor is 5.5.5.5, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 5.5.5.5
  BGP state = Established, up for 00:07:58
  Last read 00:07:58, 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 Labeled-Unicast Unicast: advertised and received
  Received 47 messages, 0 notifications, 0 in queue
  Sent 49 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is 2.2.2.2
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 4, Offset 0, Mask 0x10
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: IPv6 Labeled-Unicast Unicast
  BGP table version 6, neighbor version 6
  Index 4, Offset 0, Mask 0x10
  AF-dependant capabilities:
    Add-Path Send Capability : advertised
    Add-Path Receive Capability : received
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  2 announced prefixes

 Connections established 2; dropped 1
Local host: 2.2.2.2, Local port: 52433
Foreign host: 5.5.5.5, Foreign port: 179
Nexthop: 2.2.2.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
```

```
Last Reset: 00:08:03, due to BGP Notification sent
Notification Error Message: (Cease/Other Configuration Change.)
```

IPV6 LABELED-UNICAST Prefix Advertised to All 6pe Peers With All Addpaths:

```
RR#show bgp ipv6 labeled-unicast
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

*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
*ail                     ::ffff:4.4.4.4                      0             100          0       300 i
*ail                     ::ffff:5.5.5.5                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 3
RR#

RR#show bgp ipv6 labeled-unicast2222::2

  300, (Received from a RR-client)
    ::ffff:3.3.3.3 (metric 11) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       1.1.1.1       4.4.4.4       5.5.5.5
      Last update: Fri Mar 15 04:09:48 2019

  300, (Received from a RR-client)
    ::ffff:4.4.4.4 (metric 11) from 4.4.4.4 (4.4.4.4)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, add-path

      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       1.1.1.1       3.3.3.3       5.5.5.5
      Last update: Fri Mar 15 04:09:48 2019

  300, (Received from a RR-client)
    ::ffff:5.5.5.5 (metric 11) from 5.5.5.5 (5.5.5.5)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, add-path

      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       1.1.1.1       3.3.3.3       4.4.4.4
      Last update: Fri Mar 15 04:09:48 2019
```

PE1

```
PE1#show bgp ipv6 labeled-unicast
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
*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
```

```
* il                    ::ffff:4.4.4.4                    0             100          0       300 i
* il                    ::ffff:5.5.5.5                    0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 3
PE1#
```

PE2

```
PE2#show bgp ipv6 labeled-unicast
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
*>il 2222::2/128         ::ffff:4.4.4.4                    0             100          0       300 i
* il                    ::ffff:5.5.5.5                    0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 2
PE2#
PE2#show bgp ipv6 labeled-unicast 2222::2
  300
    2001::2(fe80::5054:ff:fe18:619a) from 2001::2 (10.12.65.71)
    (fe80::5054:ff:fe18:619a)
      Origin IGP, metric 0, localpref 100, valid, external, best

      Not advertised to any peer
      Last update: Fri Mar 15 03:57:09 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019


  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019

  300
    ::ffff:5.5.5.5 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 5.5.5.5, Cluster list: 2.2.2.2
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019
```

PE3

```
PE3#show bgp ipv6 labeled-unicast
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
*>il 2222::2/128         ::ffff:3.3.3.3                       0             100          0       300 i
* il                     ::ffff:5.5.5.5                       0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 2
PE3#
PE3#show bgp ipv6 labeled-unicast 2222::2
BGP routing table entry for 2222::2
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:09:56 2019

  300
    ::ffff:5.5.5.5 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 5.5.5.5, Cluster list: 2.2.2.2
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019
```

PE4

```
PE4#show bgp ipv6 labeled-unicast
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
*>il 2222::2/128         ::ffff:3.3.3.3                       0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 1
PE4#
PE4#show bgp ipv6 labeled-unicast 2222::2
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:09:57 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:09:57 2019
```

© 2026 IP Infusion Inc. Proprietary

CE1

```
CE1#show bgp ipv6
BGP table version is 2, local router ID is 10.12.65.70
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
*>   2222::2/128         1001::2(fe80::5054:ff:feaa:791f)
                                        0             100          0       100 300 i

Total number of prefixes 1
```

BGP Addpath Selection Configuration Best 2 Under Address-family ipv6 labeled-unicast

RR

`RR(config)#router bgp 100`	Configure router bgp mode
`RR(config-router)#address-family ipv6 labeled-unicast`	Enter BGP address-family ipv6 labeled-unicast mode
`RR(config-router-af)#no bgp additional-paths select all`	Unconfigure RR to select All addpaths
`RR(config-router-af)#bgp additional-paths select best 2`	Configure RR to select best 2 addpaths
`RR(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`RR(config-router)#exit`	Exit Router BGP mode.
`RR(config)#commit`	Apply commit
`RR(config)#exit`	Exit Configure mode.

Validation

RR

```
RR#show bgp ipv6 labeled-unicast
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

*>il 2222::2/128         ::ffff:3.3.3.3                       0             100          0       300 i
*ail                     ::ffff:4.4.4.4                       0             100          0       300 i
* il                     ::ffff:5.5.5.5                       0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 3
RR#
RR#
RR#show bgp ipv6 labeled-unicast2222::2
```

```
  300, (Received from a RR-client)
    ::ffff:3.3.3.3 (metric 11) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       1.1.1.1       4.4.4.4       5.5.5.5
      Last update: Fri Mar 15 04:09:48 2019

  300, (Received from a RR-client)
    ::ffff:4.4.4.4 (metric 11) from 4.4.4.4 (4.4.4.4)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, add-path

      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       1.1.1.1       3.3.3.3       5.5.5.5
      Last update: Fri Mar 15 04:09:48 2019

  300, (Received from a RR-client)
    ::ffff:5.5.5.5 (metric 11) from 5.5.5.5 (5.5.5.5)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Not advertised to any peer
      Last update: Fri Mar 15 04:09:48 2019
```

PE1

```
PE1#show bgp ipv6 labeled-unicast
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
 (Default for Unicast)
*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 1

*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
* il                      ::ffff:4.4.4.4                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 2
PE1#
PE1#show bgp ipv6 labeled-unicast2222::2
 (Default for Unicast)
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019


  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:36:25 2019
```

```
  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:37:00 2019
```

PE2

```
PE2#show bgp ipv6 labeled-unicast
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
 (Default for Unicast)
*>   2222::2/128         2001::2(fe80::5054:ff:fe18:619a)
                                       0             100          0       300 i
* il                     ::ffff:4.4.4.4                      0             100          0       300 i
 Announced routes count = 1
 Accepted routes count = 1

*>il 2222::2/128         ::ffff:4.4.4.4                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 1
PE2#
PE2#show bgp ipv6 labeled-unicast2222::2

  300
    2001::2(fe80::5054:ff:fe18:619a) from 2001::2 (10.12.65.71)
    (fe80::5054:ff:fe18:619a)
      Origin IGP, metric 0, localpref 100, valid, external, best

      Not advertised to any peer
      Last update: Fri Mar 15 03:57:09 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:37:00 2019


  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:37:00 2019
```

PE3

```
PE3#show bgp ipv6 labeled-unicast
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

*>il 2222::2/128         ::ffff:3.3.3.3                       0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 1
PE3#
PE3#show bgp ipv6 labeled-unicast2222::2

  300
    3001::2(fe80::5054:ff:feee:e4fc) from 3001::2 (10.12.65.71)
    (fe80::5054:ff:feee:e4fc)
      Origin IGP, metric 0, localpref 100, valid, external, best

      Not advertised to any peer
      Last update: Fri Mar 15 03:56:59 2019

  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:09:56 2019


  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:36:22 2019
```

PE4

```
PE4#show bgp ipv6 labeled-unicast
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

*>il 2222::2/128         ::ffff:3.3.3.3                       0             100          0       300 i
* il                     ::ffff:4.4.4.4                       0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 2
PE4#
PE4#show bgp ipv6 labeled-unicast2222::2

  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:36:26 2019
```

```
  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:36:59 2019
```

BGP Addpath Selection Configuration Best 3 Under Address-family ipv6 labeled-unicast

RR

`RR(config)#router bgp 100`	Configure router bgp mode
`RR(config-router)#address-family ipv6 labeled-unicast`	Enter BGP address-family ipv6 labeled-unicast mode
`RR(config-router-af)#no bgp additional-paths select best`	Unconfigure RR to select Best 2 addpaths
`RR(config-router-af)#bgp additional-paths select best 3`	Configure RR to select best 3 addpaths
`RR(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`RR(config-router)#exit`	Exit Router BGP mode.
`RR(config)#commit`	Apply commit
`RR(config)#exit`	Exit Configure mode.

Validation

RR

```
RR#show bgp ipv6 labeled-unicast
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

*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
*ail                     ::ffff:4.4.4.4                      0             100          0       300 i
*ail                     ::ffff:5.5.5.5                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 3
RR#
RR#show bgp ipv6 labeled-unicast2222::2

  300, (Received from a RR-client)
    ::ffff:3.3.3.3 (metric 11) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       1.1.1.1       4.4.4.4       5.5.5.5
      Last update: Fri Mar 15 04:09:48 2019
```

```
  300, (Received from a RR-client)
    ::ffff:4.4.4.4 (metric 11) from 4.4.4.4 (4.4.4.4)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, add-path

      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       1.1.1.1       3.3.3.3       5.5.5.5
      Last update: Fri Mar 15 04:09:48 2019

  300, (Received from a RR-client)
    ::ffff:5.5.5.5 (metric 11) from 5.5.5.5 (5.5.5.5)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, add-path

      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       1.1.1.1       3.3.3.3       4.4.4.4
      Last update: Fri Mar 15 04:09:48 2019
```

PE1

```
PE1#show bgp ipv6 labeled-unicast
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

*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
* il                     ::ffff:4.4.4.4                      0             100          0       300 i
* il                     ::ffff:5.5.5.5                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 3
PE1#
PE1#show bgp ipv6 labeled-unicast2222::2

  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:27 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:58 2019

  300
    ::ffff:5.5.5.5 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 5.5.5.5, Cluster list: 2.2.2.2
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:58 2019
```

© 2026 IP Infusion Inc. Proprietary

PE2

```
PE2#show bgp ipv6 labeled-unicast
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


*>il 2222::2/128         ::ffff:4.4.4.4                      0             100          0       300 i
* il                     ::ffff:5.5.5.5                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 2
PE2#
PE2#show bgp ipv6 labeled-unicast2222::2

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:59 2019

  300
    ::ffff:5.5.5.5 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 5.5.5.5, Cluster list: 2.2.2.2
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:59 2019
```

PE3

```
PE3#show bgp ipv6 labeled-unicast
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

*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
* il                     ::ffff:5.5.5.5                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 2
PE3#
PE3#show bgp ipv6 labeled-unicast2222::2

  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:26 2019

  300
    ::ffff:5.5.5.5 (metric 12) from 2.2.2.2 (2.2.2.2)
```

```
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 5.5.5.5, Cluster list: 2.2.2.2
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:57 2019
```

PE4

```
PE4#show bgp ipv6 labeled-unicast
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

*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
* il                     ::ffff:4.4.4.4                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 2
PE4#show bgp ipv6 labeled-unicast2222::2

  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:28 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:59 2019
```

BGP Addpath Selection Configuration all At Neighbor-level under Address-family ipv6 labeled-unicast

RR

`RR(config)#router bgp 100`	Configure router bgp mode
`RR(config-router)#address-family ipv6 labeled-unicast`	Enter BGP address-family ipv6 labeled-unicast mode
`RR(config-router-af)#neighbor 1.1.1.1 advertise additional-paths all`	Configure RR to advertise All addpaths to PE1
`RR(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`RR(config-router)#exit`	Exit Router BGP mode.
`RR(config)#commit`	Apply commit
`RR(config)#exit`	Exit Configure mode.

© 2026 IP Infusion Inc. Proprietary

Validation

PE1

```
PE1#show bgp ipv6 labeled-unicast
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

*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
* il                     ::ffff:4.4.4.4                      0             100          0       300 i
* il                     ::ffff:5.5.5.5                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 3
PE1#show bgp ipv6 labeled-unicast2222::2

  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 06:12:37 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 06:12:57 2019

  300
    ::ffff:5.5.5.5 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 5.5.5.5, Cluster list: 2.2.2.2
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 06:12:57 2019
```

BGP Addpath Selection Configuration Best 2 at Neighbor-level under Address-family ipv6 labeled-unicast

RR

`RR(config)#router bgp 100`	Configure router bgp mode
`RR(config-router)#address-family ipv6 labeled-unicast`	Enter BGP address-family ipv6 labeled-unicast mode
`RR(config-router-af)#no neighbor 1.1.1.1 advertise additional-paths all`	Configure RR to advertise All addpaths to PE1
`RR(config-router-af)#neighbor 1.1.1.1 advertise additional-paths best 2`	Configure RR to advertise Best 2 addpaths to PE1

© 2026 IP Infusion Inc. Proprietary

RR(config-router-af)#exit-address-family	Exit bgp address-family mode.
RR(config-router)#exit	Exit Router BGP mode.
RR(config)#commit	Apply commit
RR(config)#exit	Exit Configure mode.

Validation

PE1

```
PE1#show bgp ipv6 labeled-unicast
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

*>il 2222::2/128         ::ffff:3.3.3.3                     0             100          0       300 i
* il                     ::ffff:4.4.4.4                     0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 2
PE1#show bgp ipv6 labeled-unicast2222::2

  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 06:19:18 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 06:19:58 2019
```

BGP Addpath Selection Configuration Best 3 At Neighbor-level under Address-family ipv6 labeled-unicast

RR

RR(config)#router bgp 100	Configure router bgp mode
RR(config-router)#address-family ipv6 labeled-unicast	Enter BGP address-family ipv6 labeled-unicast mode
RR(config-router-af)#no neighbor 1.1.1.1 advertise additional-paths best	Configure RR to advertise All addpaths to PE1
RR(config-router-af)#neighbor 1.1.1.1 advertise additional-paths best 3	Configure RR to advertise Best 3 addpaths to PE1

© 2026 IP Infusion Inc. Proprietary

RR(config-router-af)#exit-address-family	Exit bgp address-family mode.
RR(config-router)#exit	Exit Router BGP mode.
RR(config)#commit	Apply commit
RR(config)#exit	Exit Configure mode.

Validation

PE1

```
PE1#show bgp ipv6 labeled-unicast
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

*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
* il                     ::ffff:4.4.4.4                      0             100          0       300 i
* il                     ::ffff:5.5.5.5                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 3
PE1#
PE1#show bgp ipv6 labeled-unicast2222::2

  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 06:21:35 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 06:22:01 2019

  300
    ::ffff:5.5.5.5 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal

      Originator: 5.5.5.5, Cluster list: 2.2.2.2
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 06:22:01 2019
```

Disable BGP Addpath for a Specific 6pe Neighbor

RR

`RR(config)#router bgp 100`	Configure router bgp mode
`RR(config-router)#address-family ipv6 labeled-unicast`	Enter BGP address-family ipv6 labeled-unicast mode
`RR(config-router-af)#no neighbor 1.1.1.1 advertise additional-paths best`	Configure RR to advertise All addpaths to PE1
`RR(config-router-af)#neighbor 1.1.1.1 additional-paths disable`	Configure RR to Disable Addpath for the 6pe neighbor PE1
`RR(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`RR(config-router)#exit`	Exit Router BGP mode.
`RR(config)#commit`	Apply commit
`RR(config)#exit`	Exit Configure mode.

Validation

PE1

```
PE1#show bgp ipv6 labeled-unicast
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

*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 1
PE1#show bgp ipv6 labeled-unicast2222::2

  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best

      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      Not advertised to any peer
      Last update: Fri Mar 15 06:25:58 2019
```

CE1

```
CE1#show bgp ipv6 summary
BGP router identifier 10.12.65.66, 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
1001::2                  4   100   5         6       1      0      0     00:01:40           0
```

```
Total number of neighbors 1
Total number of Established sessions 1
CE1#

CE1#show bgp ipv6
BGP table version is 1, local router ID is 10.12.65.66
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
*>   1111::/64           ::                    0             100          32768    i

Total number of prefixes 1
CE1#
```

CE2

```
CE2#show bgp ipv6 summary
BGP router identifier 10.12.65.67, local AS number 300
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
2002::2                 4   100   7         8       1      0      0     00:02:44           0

Total number of neighbors 1

Total number of Established sessions 1
CE2#

CE2#show bgp ipv6
BGP table version is 1, local router ID is 10.12.65.67
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
*>   1111::/64           ::                    0             100          32768    i

Total number of prefixes 1
CE2#
```

PE1

```
PE1#show bgp ipv6
BGP table version is 2, 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
*>   1111::/64           1001::1(fe80::5054:ff:fe88:95f3)
                                      0             100          0       200 i

Total number of prefixes 1
PE1#
PE1#show mpls ilm-table
```

```
Codes: > - installed ILM, * - selected ILM, p - stale ILM
        K - CLI ILM,T - MPLS-TP, F - FRR ILM, u - FRR ILM In Use

Code  FEC                ILM-ID      In-Label    Out-Label   In-Intf   Out-
Intf     Nexthop              LSP-Type
 >    1111::/64          1           24960       N/A         N/A       N/A          127.0.0.1
    LSP_DEFAULT
PE1#
```

PE2

```
PE2#show bgp ipv6
BGP table version is 2, local router ID is 2.2.2.2
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
*>   1111::/64           2002::1(fe80::5054:ff:fe92:a520)
                                      0             100          0       300 i

Total number of prefixes 1
PE2#
PE2#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
        K - CLI ILM,T - MPLS-TP, F - FRR ILM, u - FRR ILM In Use

Code  FEC                ILM-ID      In-Label    Out-Label   In-Intf   Out-
Intf     Nexthop              LSP-Type
 >    1111::/64          1           24960       N/A         N/A       N/A          127.0.0.1
    LSP_DEFAULT
PE2#
```

RR

```
RR#show bgp ipv6 labeled-unicast all summary
BGP router identifier 3.3.3.3, local AS number 100
BGP table version is 1
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  10         9       1      0      0     00:03:36           1
2.2.2.2                  4   100   9         8       1      0      0     00:03:01           1
4.4.4.4                  4   100   9        11       1      0      0     00:03:36           0

Total number of neighbors 3
Total number of Established sessions 3
RR#


RR#show bgp ipv6 labeled-unicast all
Status codes: s suppressed, d damped, h history, a add-path, gb group-best, * valid, > best, i -
internal, S - stale
    Network             Next Hop        In Label         Out Label
*>gi 1111::/64           ::ffff:1.1.1.1                    -             24960
*agi                     ::ffff:2.2.2.2                    -             24960
RR#

RR#show bgp ipv6 labeled-unicast 1111::1/64
BGP routing table entry for 1111::/64
Paths: (2 available, best #1, table Default-IP-Routing-Table)
```

```
  Not advertised to any peer
  200
    ::ffff:1.1.1.1 (metric 11) from 1.1.1.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 100, Out-label    24960, In-label  NA , valid, internal, group-
best, best
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       4.4.4.4
      Last update: Tue Dec  3 01:50:54 2019

  300
    ::ffff:2.2.2.2 (metric 11) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, Out-label    24960, In-label  NA , valid, internal, group-
best, add-path
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       4.4.4.4
      Last update: Tue Dec  3 01:51:29 2019
RR#

RR#show bgp neighbors 4.4.4.4
BGP neighbor is 4.4.4.4, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 4.4.4.4
  BGP state = Established, up for 00:04:02
  Last read 00:04:02, 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 Labeled Unicast: advertised and received
  Received 10 messages, 0 notifications, 0 in queue
  Sent 12 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is 3.3.3.3
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 3, Offset 0, Mask 0x8
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: IPv6 Labeled-Unicast
  BGP table version 1, neighbor version 1
  Index 3, Offset 0, Mask 0x8
  AF-dependant capabilities:
    Add-Path Send Capability : advertised
    Add-Path Receive Capability : received
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 3.3.3.3, Local port: 60367
Foreign host: 4.4.4.4, Foreign port: 179
Nexthop: 3.3.3.3
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

RR#
```

PE3

```
PE3#show bgp ipv6 labeled-unicast all

Status codes: s suppressed, d damped, h history, a add-path, gb group-best, * valid, > best, i -
internal, S - stale
    Network             Next Hop        In Label         Out Label
*>i 1111::/64           ::ffff:1.1.1.1                    -             24960
* i                     ::ffff:2.2.2.2                    -             24960
PE3#

PE3#show bgp ipv6
BGP table version is 1, local router ID is 4.4.4.4
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
*>il 1111::/64           ::ffff:1.1.1.1                      0             100          0       200 i
* il                     ::ffff:2.2.2.2                      0             100          0       300 i
Total number of prefixes 1
PE3#
```

CE3

```
CE3#show bgp ipv6 summary
BGP router identifier 10.12.65.69, 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
3003::1                 4   100   3         2       1      0      0     00:00:15           1

Total number of neighbors 1

Total number of Established sessions 1
CE3#

CE3#show bgp ipv6
BGP table version is 2, local router ID is 10.12.65.69
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
*>   1111::/64           3003::1(fe80::5054:ff:fef2:c8dc)
                                          0             100          0       100 200 i

Total number of prefixes 1
CE3#
```

Additional Paths for 6VPE

This section explains how to configure the 6VPE additional path feature to enable advertising additional paths for a 6VPE prefix, over a 6VPE session.

Topology

The topology below illustrates CE1 and CE2 as customer edge routers connected to the Provider Edge routers PE1 and PE2 respectively over a non-default VRF. CE2 is advertising the IPv6 prefix 2222::2/128 over non-default VRF to PE2, PE3, and PE4 which further advertises the prefix to RR over 6VPE sessions. RR has the 6VPE prefix 2222::2/128 with 3 next-hops PE2, PE3, and PE4. In a normal scenario, RR advertises the 6VPE prefix 2222::2/128 to PE1 with only one path (the next-hop which is marked the “best”). But with 6VPE addpath configured on RR and PE1, RR can be made to advertise the 6VPE prefix with additional-paths to PE1

Figure 85. Additional Paths for 6VPE

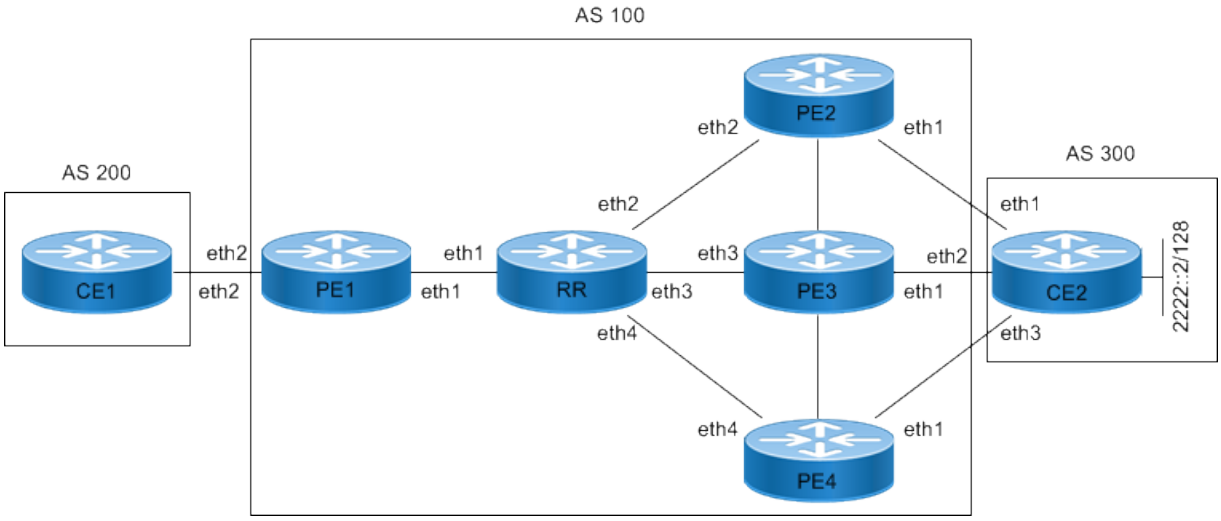

Configuration

CE1: Interface

`CE1#configure terminal`	Enter configure mode.
`CE1(config)#hostname CE1`	Configure the hostname to CE1
`CE1(config)#interface eth2`	Enter the Interface mode for interface eth2.
`CE1(config-if)#ipv6 address 1001::1/64`	Configure IPV6 address for interface eth2
`(config)#commit`	Apply commit
`(config)#exit`	Exit Router BGP and Configure mode.

CE1: BGP

`CE1(config)#router bgp 200`	Enter BGP router mode
`CE1(config-router)#neighbor 1001::2 remote-as 100`	Configure BGP ipv6 neighbor on PE1
`CE1(config-router)#address-family ipv6 unicast`	Enter the address family ipv6
`CE1(config-router-af)#neighbor 1001::2 activate`	Activate neighbor under address family ipv6
`CE1(config-router-af)#exit-address-family`	Exit address family ipv6 unicast mode
`CE1(config-router)#exit`	Exit BGP router mode

© 2026 IP Infusion Inc. Proprietary

CE1(config)#commit	Apply commit
CE1#clear bgp 1001::2	Clear BGP peer PE1

PE1: Interface

PE1#configure terminal	Enter configure mode.
PE1(config)#hostname PE1	Configure the hostname to PE1
PE1(config)#ip vrf vrf1	Configure VRF name
PE1(config-vrf)#rd 1:100	Configure Route Distinguisher value
PE1(config-vrf)#route-target both 1:200	Configure Route Target Value
PE1(config-vrf)#exit	Exit vrf mode
PE1(config)#interface eth2	Enter the Interface mode for interface eth2
PE1(config-if)#ip vrf forwarding vrf1	Configure ip vrf-forwarding on interface eth2
PE1(config-if)#ipv6 address 1001::2/64	Configure the IPv6 address for interface eth2
PE1(config-if)#exit	Exit interface mode
PE1(config)#interface lo	Enter the Interface mode for the loopback interface.
PE1(config-if)#ip address 1.1.1.1/32	Configure IPV4 address for loopback interface.
PE1(config-if)#exit	Exit interface mode
PE1(config)#interface eth1	Enter the Interface mode for interface eth1
PE1(config-if)#ip address 10.1.1.1/24	Configure the IP address for interface eth1
(config-if)#exit	Exit the interface level
(config)#commit	Apply commit
(config)#exit	Exit.

PE1: OSPF

PE1(config)#router ospf 1	Configure an OSPF instance 1
PE1(config-router)#network 10.1.1.0/24 area 0	Advertise the network towards P in Area 0
PE1(config-router)#network 1.1.1.1/32 area 0	Advertise the Loopback network in Area 0
PE1(config-router)#exit	Exit OSPF router mode
PE1(config)#commit	Apply commit
PE1(config)#exit	Exit the global mode

PE1: MPLS

PE1(config)#router ldp	Enable LDP Process
PE1(config-router)#transport-address ipv4 1.1.1.1	Configure Transport address for LDP
PE1(config-router)#exit	Exit router ldp mode
PE1(config)#interface eth1	Enter the interface mode for interface eth1
PE1(config-if)#label-switching	Enable label-switching on interface eth1

© 2026 IP Infusion Inc. Proprietary

PE1(config-if)#enable-ldp ipv4	Enable LDP for IPv4 on interface eth1
PE1(config-if)#exit	Exit the interface mode
PE1(config)#commit	Apply commit
PE1(config)#exit	Exit the global mode

PE1: BGP

PE1(config)#router bgp 100	Configure router bgp mode
PE1(config)#bgp router-id 1.1.1.1	Configure BGP router-id
PE1(config-router)#neighbor 2.2.2.2 remote-as 100	Configure BGP IPv4 neighbor on PE2
PE1(config-router)#neighbor 2.2.2.2 update-source 1.1.1.1	Configure the routing update source for RR as PE1's Loopback
PE1(config-router)#address-family vpnv6 unicast	Enter BGP address-family vpnv6 unicast mode
PE1(config-router-af)#neighbor 2.2.2.2 activate	Activate the neighbor RR for address-family vpnv6 unicast
PE1(config-router-af)#exit-address-family	Exit address-family ipv6 labeled-unicast mode
PE1(config-router)#address-family ipv6 vrf vrf1	Enter BGP address-family ipv6 vrf vrf1 mode
PE1(config-router)#neighbor 1001::1 remote-as 200	Configure BGP IPv6 neighbor on CE1
PE1(config-router-af)#exit-address-family	Exit address-family ipv6 vrf vrf1 mode
PE1(config-router)#exit	Exit router bgp mode
PE1(config)#exit	Exit configure mode
PE1(config)#commit	Apply commit
PE1#clear bgp ipv6 unicast 1001::1 vrf vrf1	Clear BGP neighbor CE1

RR: Interface

RR#configure terminal	Enter configure mode.
RR(config)#hostname RR	Configure hostname as RR
RR(config)#interface eth1	Enter interface mode for interface eth1
RR(config-if)#ip address 10.1.1.2/24	Configure IPv4 address for interface eth1
RR(config-if)#exit	Exit interface mode
RR(config)#interface eth2	Enter interface mode for interface eth2
RR(config-if)#ip address 20.1.1.2/24	Configure IPv4 address for interface eth2
RR(config-if)#exit	Exit interface mode
RR(config)#interface eth3	Enter interface mode for interface eth3
RR(config-if)#ip address 30.1.1.2/24	Configure IPv4 address for interface eth3
RR(config-if)#exit	Exit interface mode
RR(config)#interface eth4	Enter interface mode for interface eth4
RR(config-if)#ip address 40.1.1.2/24	Configure IPv4 address for interface eth4

© 2026 IP Infusion Inc. Proprietary

RR(config-if)#exit	Exit interface mode
RR(config)#interface lo	Enter Interface mode for the loopback interface.
RR(config-if)#ip address 2.2.2.2/32	Configure IPV4 address for loopback interface.
RR(config-if)#exit	Exit Interface mode.
PE1(config)#commit	Apply commit

RR: OSPF

RR(config)#router ospf 1	Configure router OSPF instance 1
RR(config-router)#network 10.1.1.0/24 area 0	Advertise the network towards PE1 in area 0
RR(config-router)#network 20.1.1.0/24 area 0	Advertise the network towards PE2 in area 0
RR(config-router)#network 30.1.1.0/24 area 0	Advertise the network towards PE3 in area 0
RR(config-router)#network 40.1.1.0/24 area 0	Advertise the network towards PE4 in area 0
RR(config-router)#network 2.2.2.2/32 area 0	Advertise the Loopback network in area 0
RR(config-router)#exit	Exit the router ospf mode
RR(config)#commit	Apply commit

RR: MPLS

RR(config)#router ldp	Enable LDP Process
RR(config-router)#transport-address ipv4 2.2.2.2	Configure Transport address for LDP
RR(config-router)#exit	Exit router ldp mode
RR(config)#interface eth1	Enter the interface mode for interface eth1
RR(config-if)#label-switching	Enable label-switching on interface eth1
RR(config-if)#enable-ldp ipv4	Enable LDP for IPv4 on interface eth1
RR(config-if)#exit	Exit interface mode
RR(config)#interface eth2	Enter the interface mode for interface eth2
RR(config-if)#label-switching	Enable label-switching on interface eth2
RR(config-if)#enable-ldp ipv4	Enable LDP for IPv4 on interface eth2
RR(config-if)#exit	Exit interface mode
RR(config)#interface eth3	Enter the interface mode for interface eth3
RR(config-if)#label-switching	Enable label-switching on interface eth3
RR(config-if)#enable-ldp ipv4	Enable LDP for IPv4 on interface eth3
RR(config-if)#exit	Exit interface mode
RR(config)#interface eth4	Enter the interface mode for interface eth4
RR(config-if)#label-switching	Enable label-switching on interface eth4
RR(config-if)#enable-ldp ipv4	Enable LDP for IPv4 on interface eth4
RR(config-if)#exit	Exit interface mode

© 2026 IP Infusion Inc. Proprietary

RR(config)#commit	Apply commit
RR(config)#exit	Exit the configure terminal mode

RR: BGP

RR(config)#router bgp 100	Configure router bgp mode
RR(config)#bgp router-id 2.2.2.2	Configure BGP router-id
RR(config-router)#neighbor 1.1.1.1 remote-as 100	Configure BGP neighbor on PE1
RR(config-router)#neighbor 1.1.1.1 update-source 2.2.2.2	Configure the routing update source for PE1 as RR's Loopback
RR(config-router)#neighbor 3.3.3.3 remote-as 100	Configure BGP neighbor on PE2
RR(config-router)#neighbor 3.3.3.3 update-source 2.2.2.2	Configure the routing update source for PE2 as RR's Loopback
RR(config-router)#neighbor 4.4.4.4 remote-as 100	Configure BGP neighbor on PE3
RR(config-router)#neighbor 4.4.4.4 update-source 2.2.2.2	Configure the routing update source for PE3 as RR's Loopback
RR(config-router)#neighbor 5.5.5.5 remote-as 100	Configure BGP neighbor on PE4
RR(config-router)#neighbor 5.5.5.5 update-source 2.2.2.2	Configure the routing update source for PE4 as RR's Loopback
RR(config-router)#address-family vpnv6 unicast	Enter BGP address-family vpnv6 unicast mode
RR(config-router-af)#neighbor 1.1.1.1 activate	Activate the neighbor PE1 for address-family vpnv6 unicast
RR(config-router-af)#neighbor 1.1.1.1 route-reflector-client	Configure PE1 as route-reflector-client
RR(config-router-af)#neighbor 3.3.3.3 activate	Activate the neighbor PE2 for address-family vpnv6 unicast
RR(config-router-af)#neighbor 3.3.3.3 route-reflector-client	Configure PE2 as route-reflector-client
RR(config-router-af)#neighbor 4.4.4.4 activate	Activate the neighbor PE3 for address-family vpnv6 unicast
RR(config-router-af)#neighbor 4.4.4.4 route-reflector-client	Configure PE3 as route-reflector-client
RR(config-router-af)#neighbor 5.5.5.5 activate	Activate the neighbor PE4 for address-family vpnv6 unicast
RR(config-router-af)#neighbor 5.5.5.5 route-reflector-client	Configure PE4 as route-reflector-client
RR(config-router-af)#exit-address-family	Exit BGP address-family vpnv6 unicast mode
RR(config-router)#exit	Exit router bgp mode
RR(config)#commit	Apply commit
RR(config)#exit	Exit configure mode

© 2026 IP Infusion Inc. Proprietary

PE2: Interface

`PE2#configure terminal`	Enter configure mode.
`PE2(config)#hostname PE2`	Configure the hostname to PE2
`PE2(config)#ip vrf vrf1`	Configure VRF name
`PE2(config-vrf)#rd 1:100`	Configure Route Distinguisher value
`PE2(config-vrf)#route-target both 1:200`	Configure Route Target Value
`PE2(config-vrf)#exit`	Exit vrf mode
`PE2(config)#interface eth1`	Enter the Interface mode for interface eth1
`PE2(config-if)#ip vrf forwarding vrf1`	Configure ip vrf-forwarding on interface eth1
`PE2(config-if)#ipv6 address 2001::1/64`	Configure the IPv6 address for interface eth1
`PE2(config-if)#exit`	Exit interface mode
`PE2(config)#interface lo`	Enter the Interface mode for the loopback interface.
`PE2(config-if)#ip address 3.3.3.3/32`	Configure IPV4 address for loopback interface.
`PE2(config-if)#exit`	Exit interface mode
`PE2(config)#interface eth2`	Enter the Interface mode for interface eth2
`PE2(config-if)#ip address 20.1.1.1/24`	Configure the IP address for interface eth2
`PE2(config-if)#exit`	Exit configure mode
`PE2(config)#commit`	Apply commit

PE2: OSPF

`PE2(config)#router ospf 1`	Configure an OSPF instance 1
`PE2(config-router)#network 20.1.1.0/24 area 0`	Advertise the network towards RR in Area 0
`PE2(config-router)#network 3.3.3.3/32 area 0`	Advertise the Loopback network in Area 0
`PE2(config-router)#exit`	Exit OSPF router mode
`PE2(config)#commit`	Apply commit

PE2: MPLS

`PE2(config)#router ldp`	Enable LDP Process
`PE2(config-router)#transport-address ipv4 3.3.3.3`	Configure Transport address for LDP
`PE2(config-router)#exit`	Exit router ldp mode
`PE2(config)#interface eth2`	Enter the interface mode for interface eth2
`PE2(config-if)#label-switching`	Enable label-switching on interface eth2
`PE2(config-if)#enable-ldp ipv4`	Enable LDP for IPv4 on interface eth2
`PE2(config-if)#exit`	Exit interface mode
`PE2(config)#commit`	Apply commit

© 2026 IP Infusion Inc. Proprietary

PE2: BGP

PE2(config)#router bgp 100	Configure router bgp mode
PE2(config)#bgp router-id 3.3.3.3	Configure BGP router-id
PE2(config-router)#neighbor 2.2.2.2 remote-as 100	Configure BGP IPv4 neighbor on RR
PE2(config-router)#neighbor 2.2.2.2 update-source 3.3.3.3	Configure the routing update source for RR as PE2's Loopback
PE2(config-router)#address-family vpnv6 unicast	Enter BGP address-family vpnv6 unicast mode
PE2(config-router-af)#neighbor 2.2.2.2 activate	Activate the neighbor RR for address-family vpnv6 unicast
PE2(config-router-af)#exit-address-family	Exit address-family ipv6 labeled-unicast mode
PE2(config-router)#address-family ipv6 vrf vrf1	Enter BGP address-family ipv6 vrf vrf1 mode
PE2(config-router)#neighbor 2001::2 remote-as 300	Configure BGP IPv6 neighbor on CE2
PE2(config-router-af)#exit-address-family	Exit address-family ipv6 vrf vrf1 mode
PE2(config-router)#exit	Exit router bgp mode
PE2(config)#commit	Apply commit
PE2(config)#exit	Exit configure mode
PE2#clear bgp ipv6 unicast 2001::2 vrf vrf1	Clear BGP neighbor CE2

PE3: Interface

PE3#configure terminal	Enter configure mode.
PE3(config)#hostname PE3	Configure the hostname to PE3
PE3(config)#ip vrf vrf1	Configure VRF name
PE3(config-vrf)#rd 1:100	Configure Route Distinguisher value
PE3(config-vrf)#route-target both 1:200	Configure Route Target Value
PE3(config-vrf)#exit	Exit vrf mode
PE3(config)#interface eth1	Enter the Interface mode for interface eth1
PE3(config-if)#ip vrf forwarding vrf1	Configure ip vrf-forwarding on interface eth1
PE3(config-if)#ipv6 address 3001::1/64	Configure the IPv6 address for interface eth1
PE3(config-if)#exit	Exit interface mode
PE3(config)#interface lo	Enter the Interface mode for the loopback interface.
PE3(config-if)#ip address 4.4.4.4/32	Configure IPV4 address for loopback interface.
PE3(config-if)#exit	Exit interface mode
PE3(config)#interface eth3	Enter the Interface mode for interface eth3
PE3(config-if)#ip address 30.1.1.1/24	Configure the IP address for interface eth3
PE3(config-if)#exit	Exit interface mode
PE3(config)#commit	Apply commit

© 2026 IP Infusion Inc. Proprietary

PE3: OSPF

`PE3(config)#router ospf 1`	Configure an OSPF instance 1
`PE3(config-router)#network 30.1.1.0/24 area 0`	Advertise the network towards RR in Area 0
`PE3(config-router)#network 4.4.4.4/32 area 0`	Advertise the Loopback network in Area 0
`PE3(config-router)#exit`	Exit OSPF router mode
`PE3(config)#commit`	Apply commit

PE3: MPLS

`PE3(config)#router ldp`	Enable LDP Process
`PE3(config-router)#transport-address ipv4 4.4.4.4`	Configure Transport address for LDP
`PE3(config-router)#exit`	Exit router ldp mode
`PE3(config)#interface eth3`	Enter the interface mode for interface eth3
`PE3(config-if)#label-switching`	Enable label-switching on interface eth3
`PE3(config-if)#enable-ldp ipv4`	Enable LDP for IPv4 on interface eth3
`PE3(config-if)#exit`	Exit interface mode
`PE3(config)#commit`	Apply commit

PE3: BGP

`PE3(config)#router bgp 100`	Configure router bgp mode
`PE3(config)#bgp router-id 4.4.4.4`	Configure BGP router-id
`PE3(config-router)#neighbor 2.2.2.2 remote-as 100`	Configure BGP IPv4 neighbor on RR
`PE3(config-router)#neighbor 2.2.2.2 update-source 4.4.4.4`	Configure the routing update source for RR as PE3's Loopback
`PE3(config-router)#address-family vpnv6 unicast`	Enter BGP address-family vpnv6 unicast mode
`PE3(config-router-af)#neighbor 2.2.2.2 activate`	Activate the neighbor RR for address-family vpnv6 unicast
`PE3(config-router-af)#exit-address-family`	Exit address-family ipv6 labeled-unicast mode
`PE3(config-router)#address-family ipv6 vrf vrf1`	Enter BGP address-family ipv6 vrf vrf1 mode
`PE3(config-router)#neighbor 3001::2 remote-as 300`	Configure BGP IPv6 neighbor on CE2
`PE3(config-router-af)#exit-address-family`	Exit address-family ipv6 vrf vrf1 mode
`PE3(config-router)#exit`	Exit router bgp mode
`PE3(config)#commit`	Apply commit
`PE3(config)#exit`	Exit configure mode
`PE3#clear bgp ipv6 unicast * vrf vrf1`	Reset IPv6 BGP connection for all VRF addresses.

© 2026 IP Infusion Inc. Proprietary

PE4: Interface

`PE4#configure terminal`	Enter configure mode.
`PE4(config)#hostname PE4`	Configure the hostname to PE4
`PE4(config)#ip vrf vrf1`	Configure VRF name
`PE4(config-vrf)#rd 1:100`	Configure Route Distinguisher value
`PE4(config-vrf)#route-target both 1:200`	Configure Route Target Vlaue
`PE4(config-vrf)#exit`	Exit vrf mode
`PE4(config)#interface eth1`	Enter the Interface mode for interface eth1
`PE4(config-if)#ip vrf forwarding vrf1`	Configure ip vrf-forwarding on interface eth1
`PE4(config-if)#ipv6 address 4001::1/64`	Configure the IPv6 address for interface eth1
`PE4(config-if)#exit`	Exit interface mode
`PE4(config)#interface lo`	Enter the Interface mode for the loopback interface.
`PE4(config-if)#ip address 5.5.5.5/32`	Configure IPV4 address for loopback interface.
`PE4(config-if)#exit`	Exit interface mode
`PE4(config)#interface eth4`	Enter the Interface mode for interface eth4
`PE4(config-if)#ip address 40.1.1.1/24`	Configure the IP address for interface eth4
`PE4(config-if)#exit`	Exit interface mode
`PE4(config)#commit`	Apply commit

PE4: OSPF

`PE4(config)#router ospf 1`	Configure an OSPF instance 1
`PE4(config-router)#network 40.1.1.0/24 area 0`	Advertise the network towards RR in Area 0
`PE4(config-router)#network 5.5.5.5/32 area 0`	Advertise the Loopback network in Area 0
`PE4(config-router)#exit`	Exit OSPF router mode
`PE4(config)#commit`	Apply commit

PE4: MPLS

`PE4(config)#router ldp`	Enable LDP Process
`PE4(config-router)#transport-address ipv4 5.5.5.5`	Configure Transport address for LDP
`PE4(config-router)#exit`	Exit router ldp mode
`PE4(config)#interface eth4`	Enter the interface mode for interface eth4
`PE4(config-if)#label-switching`	Enable label-switching on interface eth4
`PE4(config-if)#enable-ldp ipv4`	Enable LDP for IPv4 on interface eth4
`PE4(config-if)#exit`	Exit interface mode
`PE4(config)#commit`	Apply commit

© 2026 IP Infusion Inc. Proprietary

PE4: BGP

PE4(config)#router bgp 100	Configure router bgp mode
PE4(config)#bgp router-id 5.5.5.5	Configure BGP router-id
PE4(config-router)#neighbor 2.2.2.2 remote-as 100	Configure BGP IPv4 neighbor on RR
PE4(config-router)#neighbor 2.2.2.2 update-source 5.5.5.5	Configure the routing update source for RR as PE4's Loopback
PE4(config-router)#address-family vpnv6 unicast	Enter BGP address-family vpnv6 unicast mode
PE4(config-router-af)#neighbor 2.2.2.2 activate	Activate the neighbor RR for address-family vpnv6 unicast
PE4(config-router-af)#exit-address-family	Exit address-family ipv6 labeled-unicast mode
PE4(config-router)#address-family ipv6 vrf vrf1	Enter BGP address-family ipv6 vrf vrf1 mode
PE4(config-router)#neighbor 4001::2 remote-as 300	Configure BGP IPv6 neighbor on CE2
PE4(config-router-af)#exit-address-family	Exit address-family ipv6 vrf vrf1 mode
PE4(config-router)#exit	Exit router bgp mode
PE4(config)#commit	Apply commit
PE4(config)#exit	Exit configure mode
PE4#clear bgp ipv6 unicast * vrf vrf1	Reset IPv6 BGP connection for all VRF addresses.

CE2: Interface

CE2#configure terminal	Enter configure mode.
CE2(config)#hostname CE2	Configure the hostname to CE2
CE2(config)#interface eth1	Enter the Interface mode for interface eth1
CE2(config-if)#ipv6 address 2001::2/64	Configure IPV6 address for interface eth1
CE2(config-if)#exit	Exit Interface mode.
CE2(config)#interface eth2	Enter the Interface mode for interface eth2
CE2(config-if)#ipv6 address 3001::2/64	Configure IPV6 address for interface eth2
CE2(config-if)#exit	Exit Interface mode.
CE2(config)#interface eth3	Enter the Interface mode for interface eth3
CE2(config-if)#ipv6 address 4001::2/64	Configure IPV6 address for interface eth3
CE2(config-if)#exit	Exit Interface mode.
CE2(config)#interface lo	Enter the Interface mode for interface Loopback
CE2(config-if)#ipv6 address 2222::2/128	Configure IPV6 address for interface Loopback
CE2(config-if)#exit	Exit Interface mode.
CE2(config)#commit	Apply commit

© 2026 IP Infusion Inc. Proprietary

CE2: BGP

CE2(config)#router bgp 300	Enter BGP router mode
CE2(config-router)#neighbor 2001::1 remote-as 100	Configure BGP ipv6 neighbor on PE2
CE2(config-router)#neighbor 3001::1 remote-as 100	Configure BGP ipv6 neighbor on PE3
CE2(config-router)#neighbor 4001::1 remote-as 100	Configure BGP ipv6 neighbor on PE4
CE2(config-router)#address-family ipv6 unicast	Enter the address family ipv6
CE2(config-router-af)#neighbor 2001::1 activate	Activate neighbor under address family ipv6
CE2(config-router-af)#neighbor 3001::1 activate	Activate neighbor under address family ipv6
CE2(config-router-af)#neighbor 4001::1 activate	Activate neighbor under address family ipv6
CE2(config-router-af)#network 2222::2/128	Advertise the IPv6 prefix of Loopback into BGP
CE2(config-router-af)#exit-address-family	Exit address family ipv6 unicast mode
CE2(config-router)#exit	Exit BGP router mode
CE2(config)#commit	Apply commit
CE2(config)#exit	Exit configure mode

BGP Addpath Capability for VPNV6 Unicast Address-family

RR

RR(config)#router bgp 100	Configure router bgp mode
RR(config-router)#address-family vpnv6 unicast	Enter BGP address-family vpnv6 unicast mode
RR(config-router-af)#bgp additional-paths select all	Configure RR to select all addpaths
RR(config-router-af)#bgp additional-paths send	Configure BGP addpath send capability under address-family vpnv6 unicast
(config-router-af)#exit-address-family	Exit bgp address-family mode.
(config-router)#exit	Exit Router BGP mode.
(config)#commit	Apply commit
(config)#exit	Exit Configure mode.

PE1

PE1(config)#router bgp 100	Configure router bgp mode
PE1(config-router)#address-family vpnv6 unicast	Enter BGP address-family vpnv6 unicast mode
PE1(config-router-af)#neighbor 2.2.2.2 additional-paths receive	Configure BGP addpath receive capability for the 6VPE-neighbor RR under address-family vpnv6 unicast
(config-router-af)#exit-address-family	Exit bgp address-family mode.
(config-router)#exit	Exit Router BGP mode.
(config)#commit	Apply commit
(config)#exit	Exit Configure mode.

© 2026 IP Infusion Inc. Proprietary

PE2

`PE2(config)#router bgp 100`	Configure router bgp mode
`PE2(config-router)#address-family vpnv6 unicast`	Enter BGP address-family vpnv6 unicast mode
`PE2(config-router-af)#bgp additional-paths receive`	Configure BGP addpath receive capability under address-family vpnv6 unicast
`(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`(config-router)#exit`	Exit Router BGP mode.
`(config)#commit`	Apply commit
`(config)#exit`	Exit Configure mode.

PE3

`PE3(config)#router bgp 100`	Configure router bgp mode
`PE3(config-router)#address-family vpnv6 unicast`	Enter BGP address-family vpnv6 unicast mode
`PE3(config-router-af)#bgp additional-paths receive`	Configure BGP addpath receive capability under address-family vpnv6 unicast
`PE3(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`PE3(config-router)#exit`	Exit Router BGP mode.
`PE3(config)#commit`	Apply commit
`PE3(config)#exit`	Exit Configure mode.

PE4

`PE4(config)#router bgp 100`	Configure router bgp mode
`PE4(config-router)#address-family vpnv6 unicast`	Enter BGP address-family vpnv6 unicast mode
`PE4(config-router-af)#bgp additional-paths receive`	Configure BGP addpath receive capability under address-family vpnv6 unicast
`PE4(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`PE4(config-router)#exit`	Exit Router BGP mode.
`PE4(config)#commit`	Apply commit
`PE4(config)#exit`	Exit Configure mode.

Validation at RR

Add-Path Send Capability Advertised to and Received From 6VPE Peers

```
RR#show bgp neighbors 1.1.1.1
BGP neighbor is 1.1.1.1, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 1.1.1.1
  BGP state = Established, up for 00:03:03
  Last read 00:03:03, 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
    Address family VPNv6 Unicast: advertised and received
```

```
  Received 35 messages, 0 notifications, 0 in queue
  Sent 41 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is 2.2.2.2
 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)
  0 accepted prefixes
  0 announced prefixes

 For address family: VPNv6 Unicast
  BGP table version 6, neighbor version 6
  Index 1, Offset 0, Mask 0x2
  AF-dependant capabilities:
    Add-Path Send Capability : advertised
    Add-Path Receive Capability : received
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  3 announced prefixes

 Connections established 2; dropped 1
Local host: 2.2.2.2, Local port: 57154
Foreign host: 1.1.1.1, Foreign port: 179
Nexthop: 2.2.2.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:03:08, due to BGP Notification sent
Notification Error Message: (Cease/Other Configuration Change.)
RR#


RR#show bgp neighbors 3.3.3.3
BGP neighbor is 3.3.3.3, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 3.3.3.3
  BGP state = Established, up for 00:05:20
  Last read 00:05:20, 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 VPNv6 Unicast: advertised and received
  Received 41 messages, 0 notifications, 0 in queue
  Sent 44 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is 2.2.2.2
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: VPNv6 Unicast
  BGP table version 6, neighbor version 6
  Index 2, Offset 0, Mask 0x4
  AF-dependant capabilities:
    Add-Path Send Capability : advertised
    Add-Path Receive Capability : received
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  2 announced prefixes
```

```
 Connections established 2; dropped 1
Local host: 2.2.2.2, Local port: 179
Foreign host: 3.3.3.3, Foreign port: 32841
Nexthop: 2.2.2.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:05:25, due to BGP Notification sent
Notification Error Message: (Cease/Other Configuration Change.)
RR#


RR#show bgp neighbors 4.4.4.4
BGP neighbor is 4.4.4.4, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 4.4.4.4
  BGP state = Established, up for 00:05:47
  Last read 00:05:47, 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 VPNv6 Unicast: advertised and received
  Received 42 messages, 0 notifications, 0 in queue
  Sent 46 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is 2.2.2.2
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 3, Offset 0, Mask 0x8
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: VPNv6 Unicast
  BGP table version 6, neighbor version 6
  Index 3, Offset 0, Mask 0x8
  AF-dependant capabilities:
    Add-Path Send Capability : advertised
    Add-Path Receive Capability : received
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  2 announced prefixes

 Connections established 2; dropped 1
Local host: 2.2.2.2, Local port: 179
Foreign host: 4.4.4.4, Foreign port: 42015
Nexthop: 2.2.2.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:05:47, due to BGP Notification received
Notification Error Message: (Cease/Administratively Reset.)
RR#


RR#show bgp neighbors 5.5.5.5
BGP neighbor is 5.5.5.5, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 5.5.5.5
  BGP state = Established, up for 00:07:58
  Last read 00:07:58, 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 VPNv6 Unicast: advertised and received
  Received 47 messages, 0 notifications, 0 in queue
```

```
  Sent 49 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is 2.2.2.2
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 4, Offset 0, Mask 0x10
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: VPNv6 Unicast
  BGP table version 6, neighbor version 6
  Index 4, Offset 0, Mask 0x10
  AF-dependant capabilities:
    Add-Path Send Capability : advertised
    Add-Path Receive Capability : received
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  2 announced prefixes

 Connections established 2; dropped 1
Local host: 2.2.2.2, Local port: 52433
Foreign host: 5.5.5.5, Foreign port: 179
Nexthop: 2.2.2.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:08:03, due to BGP Notification sent
Notification Error Message: (Cease/Other Configuration Change.)
```

VPNV6 Prefix Advertised to All 6VPE Peers With All Addpaths

```
RR#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:3.3.3.3                     0             100          0       300 i
*ail                     ::ffff:4.4.4.4                     0             100          0       300 i
*ail                     ::ffff:5.5.5.5                     0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 3
RR#

RR#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100
  300, (Received from a RR-client)
    ::ffff:3.3.3.3 (metric 11) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       1.1.1.1       4.4.4.4       5.5.5.5
      Last update: Fri Mar 15 04:09:48 2019

  300, (Received from a RR-client)
    ::ffff:4.4.4.4 (metric 11) from 4.4.4.4 (4.4.4.4)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, add-path
      Extended Community: RT:1:200
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       1.1.1.1       3.3.3.3       5.5.5.5
```

```
      Last update: Fri Mar 15 04:09:48 2019

  300, (Received from a RR-client)
    ::ffff:5.5.5.5 (metric 11) from 5.5.5.5 (5.5.5.5)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, add-path
      Extended Community: RT:1:200
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       1.1.1.1       3.3.3.3       4.4.4.4
      Last update: Fri Mar 15 04:09:48 2019
```

Validation at PE1

```
PE1#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100 (Default for VRF vrf1)
*>il 2222::2/128         ::ffff:3.3.3.3                       0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:3.3.3.3                       0             100          0       300 i
* il                     ::ffff:4.4.4.4                       0             100          0       300 i
* il                     ::ffff:5.5.5.5                       0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 3
PE1#
PE1#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100 (Default for VRF vrf1)
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019

Route Distinguisher: 1:100
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019

  300
    ::ffff:5.5.5.5 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 5.5.5.5, Cluster list: 2.2.2.2
      rx path_id: 2      tx path_id: -1
```

```
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019
```

Validation at PE2

```
PE2#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100 (Default for VRF vrf1)
*>   2222::2/128         2001::2(fe80::5054:ff:fe18:619a)
                                      0             100          0       300 i
* il                     ::ffff:4.4.4.4                     0             100          0       300 i
 Announced routes count = 1
 Accepted routes count = 1
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:4.4.4.4                     0             100          0       300 i
* il                     ::ffff:5.5.5.5                     0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 2
PE2#
PE2#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100 (Default for VRF vrf1)
  300
    2001::2(fe80::5054:ff:fe18:619a) from 2001::2 (10.12.65.71)
    (fe80::5054:ff:fe18:619a)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Extended Community: RT:1:200
      Not advertised to any peer
      Last update: Fri Mar 15 03:57:09 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019

Route Distinguisher: 1:100
  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019

  300
    ::ffff:5.5.5.5 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 5.5.5.5, Cluster list: 2.2.2.2
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019
```

Validation at PE3

```
PE3#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100 (Default for VRF vrf1)
*>   2222::2/128         3001::2(fe80::5054:ff:feee:e4fc)
                                       0             100          0       300 i
* il                     ::ffff:3.3.3.3                      0             100          0       300 i
 Announced routes count = 1
 Accepted routes count = 1
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
* il                     ::ffff:5.5.5.5                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 2
PE3#
PE3#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100 (Default for VRF vrf1)
  300
    3001::2(fe80::5054:ff:feee:e4fc) from 3001::2 (10.12.65.71)
    (fe80::5054:ff:feee:e4fc)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Extended Community: RT:1:200
      Not advertised to any peer
      Last update: Fri Mar 15 03:56:59 2019

  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:09:56 2019

Route Distinguisher: 1:100
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:09:56 2019

  300
    ::ffff:5.5.5.5 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 5.5.5.5, Cluster list: 2.2.2.2
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019
```

Validation at PE4

```
PE4#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100 (Default for VRF vrf1)
*>   2222::2/128         4001::2(fe80::5054:ff:fe2e:4cb9)
                                       0             100          0       300 i
```

```
* il                    ::ffff:3.3.3.3                     0             100          0       300 i
 Announced routes count = 1
 Accepted routes count = 1
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:3.3.3.3                     0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 1
PE4#
PE4#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100 (Default for VRF vrf1)
  300
    4001::2(fe80::5054:ff:fe2e:4cb9) from 4001::2 (10.12.65.71)
    (fe80::5054:ff:fe2e:4cb9)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Extended Community: RT:1:200
      Not advertised to any peer
      Last update: Fri Mar 15 03:56:38 2019

  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:09:57 2019

Route Distinguisher: 1:100
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:09:57 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:09:57 2019
```

Validation at CE1

```
CE1#show bgp ipv6
BGP table version is 2, local router ID is 10.12.65.70
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
*>   2222::2/128         1001::2(fe80::5054:ff:feaa:791f)
                                        0             100          0       100 300 i

Total number of prefixes 1
```

BGP Addpath Selection Configuration Best 2 Under Address-family VPNV6 Unicast

RR

`RR(config)#router bgp 100`	Configure router bgp mode
`RR(config-router)#address-family vpnv6 unicast`	Enter BGP address-family vpnv6 unicast mode
`RR(config-router-af)#no bgp additional-paths select all`	Unconfigure RR to select All addpaths
`RR(config-router-af)#bgp additional-paths select best 2`	Configure RR to select best 2 addpaths
`(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`(config-router)#exit`	Exit Router BGP mode.
`(config)#commit`	Apply commit
`(config)#exit`	Exit Configure mode.

Validation at RR

```
RR#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
*ail                     ::ffff:4.4.4.4                      0             100          0       300 i
* il                     ::ffff:5.5.5.5                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 3
RR#
RR#
RR#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100
  300, (Received from a RR-client)
    ::ffff:3.3.3.3 (metric 11) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       1.1.1.1       4.4.4.4       5.5.5.5
      Last update: Fri Mar 15 04:09:48 2019

  300, (Received from a RR-client)
    ::ffff:4.4.4.4 (metric 11) from 4.4.4.4 (4.4.4.4)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, add-path
      Extended Community: RT:1:200
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       1.1.1.1       3.3.3.3       5.5.5.5
      Last update: Fri Mar 15 04:09:48 2019

  300, (Received from a RR-client)
    ::ffff:5.5.5.5 (metric 11) from 5.5.5.5 (5.5.5.5)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Not advertised to any peer
      Last update: Fri Mar 15 04:09:48 2019
```

Validation at PE1

```
PE1#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100 (Default for VRF vrf1)
*>il 2222::2/128         ::ffff:3.3.3.3                       0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:3.3.3.3                       0             100          0       300 i
* il                     ::ffff:4.4.4.4                       0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 2
PE1#
PE1#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100 (Default for VRF vrf1)
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019

Route Distinguisher: 1:100
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:36:25 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:37:00 2019
```

Validation at PE2

```
PE2#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100 (Default for VRF vrf1)
*>   2222::2/128         2001::2(fe80::5054:ff:fe18:619a)
                                  0             100          0       300 i
* il                     ::ffff:4.4.4.4                       0             100          0       300 i
 Announced routes count = 1
 Accepted routes count = 1
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:4.4.4.4                       0             100          0       300 i
 Announced routes count = 0
```

```
 Accepted routes count = 1
PE2#
PE2#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100 (Default for VRF vrf1)
  300
    2001::2(fe80::5054:ff:fe18:619a) from 2001::2 (10.12.65.71)
    (fe80::5054:ff:fe18:619a)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Extended Community: RT:1:200
      Not advertised to any peer
      Last update: Fri Mar 15 03:57:09 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:37:00 2019

Route Distinguisher: 1:100
  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:37:00 2019
```

Validation at PE3

```
PE3#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100 (Default for VRF vrf1)
*>   2222::2/128         3001::2(fe80::5054:ff:feee:e4fc)
                                    0             100          0       300 i
* il                     ::ffff:3.3.3.3                      0             100          0       300 i
 Announced routes count = 1
 Accepted routes count = 1
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 1
PE3#
PE3#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100 (Default for VRF vrf1)
  300
    3001::2(fe80::5054:ff:feee:e4fc) from 3001::2 (10.12.65.71)
    (fe80::5054:ff:feee:e4fc)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Extended Community: RT:1:200
      Not advertised to any peer
      Last update: Fri Mar 15 03:56:59 2019

  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
```

```
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:09:56 2019

Route Distinguisher: 1:100
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:36:22 2019
```

Validation at PE4

```
PE4#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100 (Default for VRF vrf1)
*>   2222::2/128         4001::2(fe80::5054:ff:fe2e:4cb9)
                                       0             100          0       300 i
* il                     ::ffff:3.3.3.3                     0             100          0       300 i
 Announced routes count = 1
 Accepted routes count = 1
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:3.3.3.3                     0             100          0       300 i
* il                     ::ffff:4.4.4.4                     0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 2
PE4#
PE4#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100 (Default for VRF vrf1)
  300
    4001::2(fe80::5054:ff:fe2e:4cb9) from 4001::2 (10.12.65.71)
    (fe80::5054:ff:fe2e:4cb9)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Extended Community: RT:1:200
      Not advertised to any peer
      Last update: Fri Mar 15 03:56:38 2019

  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:09:57 2019

Route Distinguisher: 1:100
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:36:26 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
```

```
      Extended Community: RT:1:200
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:36:59 2019
```

BGP Addpath Selection Configuration Best 3 Under Address-family VPNV6 Unicast

RR

`RR(config)#router bgp 100`	Configure router bgp mode
`RR(config-router)#address-family vpnv6 unicast`	Enter BGP address-family vpnv6 unicast mode
`RR(config-router-af)#no bgp additional-paths select best`	Unconfigure RR to select Best 2 addpaths
`RR(config-router-af)#bgp additional-paths select best 3`	Configure RR to select best 3 addpaths
`RR(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`RR(config-router)#exit`	Exit Router BGP mode.
`RR(config)#commit`	Apply commit
`RR(config)#exit`	Exit Configure mode.

Validation at RR

```
RR#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:3.3.3.3                     0             100          0       300 i
*ail                     ::ffff:4.4.4.4                     0             100          0       300 i
*ail                     ::ffff:5.5.5.5                     0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 3
RR#
RR#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100
  300, (Received from a RR-client)
    ::ffff:3.3.3.3 (metric 11) from 3.3.3.3 (3.3.3.3)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       1.1.1.1       4.4.4.4       5.5.5.5
      Last update: Fri Mar 15 04:09:48 2019

  300, (Received from a RR-client)
    ::ffff:4.4.4.4 (metric 11) from 4.4.4.4 (4.4.4.4)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, add-path
      Extended Community: RT:1:200
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       1.1.1.1       3.3.3.3       5.5.5.5
      Last update: Fri Mar 15 04:09:48 2019

  300, (Received from a RR-client)
    ::ffff:5.5.5.5 (metric 11) from 5.5.5.5 (5.5.5.5)
```

```
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, add-path
      Extended Community: RT:1:200
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       1.1.1.1       3.3.3.3       4.4.4.4
      Last update: Fri Mar 15 04:09:48 2019
```

Validation at PE1

```
PE1#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100 (Default for VRF vrf1)
*>il 2222::2/128         ::ffff:3.3.3.3                       0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:3.3.3.3                       0             100          0       300 i
* il                     ::ffff:4.4.4.4                       0             100          0       300 i
* il                     ::ffff:5.5.5.5                       0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 3
PE1#
PE1#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100 (Default for VRF vrf1)
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019

Route Distinguisher: 1:100
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:27 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:58 2019

  300
    ::ffff:5.5.5.5 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 5.5.5.5, Cluster list: 2.2.2.2
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:58 2019
```

Validation at PE2

```
PE2#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100 (Default for VRF vrf1)
*>   2222::2/128         2001::2(fe80::5054:ff:fe18:619a)
                                        0             100          0       300 i
* il                     ::ffff:4.4.4.4                      0             100          0       300 i
 Announced routes count = 1
 Accepted routes count = 1
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:4.4.4.4                      0             100          0       300 i
* il                     ::ffff:5.5.5.5                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 2
PE2#
PE2#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100 (Default for VRF vrf1)
  300
    2001::2(fe80::5054:ff:fe18:619a) from 2001::2 (10.12.65.71)
    (fe80::5054:ff:fe18:619a)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Extended Community: RT:1:200
      Not advertised to any peer
      Last update: Fri Mar 15 03:57:09 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:59 2019

Route Distinguisher: 1:100
  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:59 2019

  300
    ::ffff:5.5.5.5 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 5.5.5.5, Cluster list: 2.2.2.2
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:59 2019
```

Validation at PE3

```
PE3#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100 (Default for VRF vrf1)
*>   2222::2/128         3001::2(fe80::5054:ff:feee:e4fc)
                                       0             100          0       300 i
* il                     ::ffff:3.3.3.3                      0             100          0       300 i
 Announced routes count = 1
 Accepted routes count = 1
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
* il                     ::ffff:5.5.5.5                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 2
PE3#
PE3#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100 (Default for VRF vrf1)
  300
    3001::2(fe80::5054:ff:feee:e4fc) from 3001::2 (10.12.65.71)
    (fe80::5054:ff:feee:e4fc)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Extended Community: RT:1:200
      Not advertised to any peer
      Last update: Fri Mar 15 03:56:59 2019

  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:09:56 2019

Route Distinguisher: 1:100
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:26 2019

  300
    ::ffff:5.5.5.5 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 5.5.5.5, Cluster list: 2.2.2.2
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:57 2019
```

Validation at PE4

```
PE4#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100 (Default for VRF vrf1)
*>   2222::2/128         4001::2(fe80::5054:ff:fe2e:4cb9)
                                       0             100          0       300 i
* il                     ::ffff:3.3.3.3                      0             100          0       300 i
 Announced routes count = 1
 Accepted routes count = 1
Route Distinguisher: 1:100
```

```
*>il 2222::2/128         ::ffff:3.3.3.3                    0             100          0       300 i
* il                     ::ffff:4.4.4.4                    0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 2
PE4#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100 (Default for VRF vrf1)
  300
    4001::2(fe80::5054:ff:fe2e:4cb9) from 4001::2 (10.12.65.71)
    (fe80::5054:ff:fe2e:4cb9)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Extended Community: RT:1:200
      Not advertised to any peer
      Last update: Fri Mar 15 03:56:38 2019

  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:09:57 2019

Route Distinguisher: 1:100
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:28 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 05:55:59 2019
```

BGP Addpath Selection Configuration all At Neighbor-level under Address-family VPNV6 Unicast

RR

`RR(config)#router bgp 100`	Configure router bgp mode
`RR(config-router)#address-family vpnv6 unicast`	Enter BGP address-family vpnv6 unicast mode
`RR(config-router-af)#neighbor 1.1.1.1 advertise additional-paths all`	Configure RR to advertise All addpaths to PE1
`RR(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`RR(config-router)#exit`	Exit Router BGP mode.
`RR(config)#commit`	Apply commit
`RR(config)#exit`	Exit Configure mode.

Validation at PE1

```
PE1#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100 (Default for VRF vrf1)
*>il 2222::2/128         ::ffff:3.3.3.3                       0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:3.3.3.3                       0             100          0       300 i
* il                     ::ffff:4.4.4.4                       0             100          0       300 i
* il                     ::ffff:5.5.5.5                       0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 3
PE1#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100 (Default for VRF vrf1)
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019

Route Distinguisher: 1:100
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 06:12:37 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 06:12:57 2019

  300
    ::ffff:5.5.5.5 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 5.5.5.5, Cluster list: 2.2.2.2
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 06:12:57 2019
```

BGP Addpath Selection Configuration Best 2 at Neighbor-level under Address-family VPNV6 Unicast

RR

`RR(config)#router bgp 100`	Configure router bgp mode
`RR(config-router)#address-family vpnv6 unicast`	Enter BGP address-family vpnv6 unicast mode
`RR(config-router-af)#no neighbor 1.1.1.1 advertise additional-paths all`	Configure RR to advertise All addpaths to PE1

`RR(config-router-af)#neighbor 1.1.1.1 advertise additional-paths best 2`	Configure RR to advertise Best 2 addpaths to PE1
`RR(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`RR(config-router)#exit`	Exit Router BGP mode.
`RR(config)#commit`	Apply commit
`RR(config)#exit`	Exit Configure mode.

Validation at PE1

```
PE1#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100 (Default for VRF vrf1)
*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
* il                     ::ffff:4.4.4.4                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 2
PE1#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100 (Default for VRF vrf1)
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019

Route Distinguisher: 1:100
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 06:19:18 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 06:19:58 2019
```

BGP Addpath Selection Configuration Best 3 At Neighbor-level under Address-family VPNV6 Unicast

RR

`RR(config)#router bgp 100`	Configure router bgp mode

`RR(config-router)#address-family vpnv6 unicast`	Enter BGP address-family vpnv6 unicast mode
`RR(config-router-af)#no neighbor 1.1.1.1 advertise additional-paths best`	Configure RR to advertise All addpaths to PE1
`RR(config-router-af)#neighbor 1.1.1.1 advertise additional-paths best 3`	Configure RR to advertise Best 3 addpaths to PE1
`RR(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`RR(config-router)#exit`	Exit Router BGP mode.
`RR(config)#commit`	Apply commit
`RR(config)#exit`	Exit Configure mode.

Validation at PE1

```
PE1#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100 (Default for VRF vrf1)
*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
* il                     ::ffff:4.4.4.4                      0             100          0       300 i
* il                     ::ffff:5.5.5.5                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 3
PE1#
PE1#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100 (Default for VRF vrf1)
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 04:10:00 2019

Route Distinguisher: 1:100
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 06:21:35 2019

  300
    ::ffff:4.4.4.4 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 4.4.4.4, Cluster list: 2.2.2.2
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 06:22:01 2019

  300
    ::ffff:5.5.5.5 (metric 12) from 2.2.2.2 (2.2.2.2)
```

```
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal
      Extended Community: RT:1:200
      Originator: 5.5.5.5, Cluster list: 2.2.2.2
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Fri Mar 15 06:22:01 2019
```

Disable BGP Addpath for a Specific 6VPE Neighbor

RR

`RR(config)#router bgp 100`	Configure router bgp mode
`RR(config-router)#address-family vpnv6 unicast`	Enter BGP address-family vpnv6 unicast mode
`RR(config-router-af)#no neighbor 1.1.1.1 advertise additional-paths best`	Configure RR to advertise All addpaths to PE1
`RR(config-router-af)#neighbor 1.1.1.1 additional-paths disable`	Configure RR to Disable Addpath for the 6VPE neighbor PE1
`RR(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`RR(config-router)#exit`	Exit Router BGP mode.
`RR(config)#commit`	Apply commit
`RR(config)#exit`	Exit Configure mode.

Validation at PE1

```
PE1#show ip bgp vpnv6 all
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
Route Distinguisher: 1:100 (Default for VRF vrf1)
*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 1:100
*>il 2222::2/128         ::ffff:3.3.3.3                      0             100          0       300 i
 Announced routes count = 0
 Accepted routes count = 1
PE1#show ip bgp vpnv6 all 2222::2
Route Distinguisher: 1:100 (Default for VRF vrf1)
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      Not advertised to any peer
      Last update: Fri Mar 15 06:25:58 2019

Route Distinguisher: 1:100
  300
    ::ffff:3.3.3.3 (metric 12) from 2.2.2.2 (2.2.2.2)
      Origin IGP, metric 0, localpref 100, label    24960, valid, internal, best
      Extended Community: RT:1:200
      Originator: 3.3.3.3, Cluster list: 2.2.2.2
      Not advertised to any peer
      Last update: Fri Mar 15 06:25:58 2019
```

BGP Labeled Unicast IPv4 with Add Path

This section contains basic BGP labeled unicast with Add-Path configuration examples with OSPF as the IGP.

Overview

As MPLS deployments expand beyond the service provider core and edge to the access and metropolitan networks, the number of edge-to-edge label-switched paths (LSPs) in many networks is increasing substantially. This continued growth can present scaling challenges-with some networks already reaching their limit-and can also slow end-to end restoration. Dividing the network into multiple regions can alleviate these issues by limiting the total number of end-to-end LSPs, and enabling failures to be contained and restored in a single region. These regions operate separate instances of interior gateway protocol (IGP), and uses BGP Labeled Unicast (BGP-LU) to advertise route information between inter region routers. By providing connectivity and communication between regions, BPG-LU enables service providers to massively scale the number of MPLS-enabled devices on their networks.

By default, all BGP routers and Route-Reflectors propagate only their best paths over their sessions. In case they advertise any route with the same NLRI as a previously advertised route, the latest one implicitly replaces the previous advertisement, which is known as an Implicit Withdraw. The Implicit Withdraw can achieve better scaling, but at the cost of path diversity.

The use of route-reflectors thus has significant effect on redundancy by hiding alternate paths. Using full-mesh is not an option, so a mechanism is needed to allow the propagation of multiple alternate paths in an RR environment.

Topology

Figure 86. BGP Labeled unicast IPv4 with additional paths

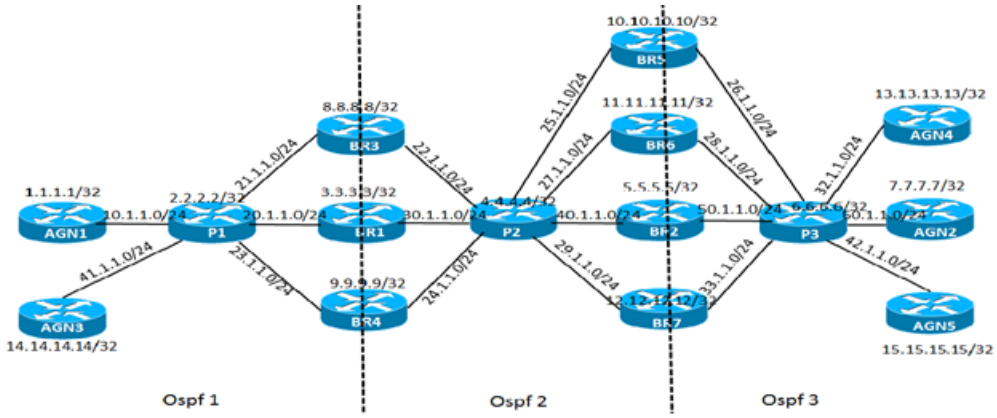

Interface Connections

AGN1(eth1) — P1(eth1)

AGN3(eth5) — P1(eth5)

P1(eth4) — BR3

P1(eth2) — BR1

P1(eth3) — BR4

BR3(eth3) — P2

BR1(eth1) — P2

BR4(eth4) — P2

P2(eth5) — BR5(eth3) — P3(eth6) — AGN4

P2(eth6) — BR6(eth4) — P3(eth2) — AGN2

P2(eth2) — BR2(eth1) — P3(eth7) — AGN5

P2(eth7) — BR7(eth5) — P3

Configurations

AGN3 Configurations

`AGN3(config)#interface lo`	Enter interface lo mode
`AGN3(config-if)#ip address 14.14.14.14/32`	Configure ip address for loopback interface
`AGN3(config-if)#exit`	Exit interface loopback mode
`AGN3(config)#router ldp`	Configure router ldp
`AGN3(config-router)#transport-address ipv4 14.14.14.14`	Configure transport address as loopback
`AGN3(config-router)#exit`	Exit the router ldp mode
`AGN3(config-if)#interface eth5`	Enter interface mode
`AGN3(config-if)#ip address 41.1.1.171/24`	Configure ip address on the interface
`AGN3(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`AGN3(config-if)#label-switching`	Enable label-switching on the interface
`AGN3(config-if)#exit`	Exit the interface mode
`AGN3(config)#router ospf 1`	Configure router ospf instance 1
`AGN3(config-router)#network 14.14.14.14/32 area 0`	Add the loopback network into ospf instance 1
`AGN3(config-router)#network 41.1.1.171/24 area 0`	Add the ip address into ospf instance 1
`AGN3(config-router)#exit`	Exit router ospf mode
`AGN3(config)#router bgp 100`	Enter router bgp mode
`AGN3(config-router)#network 14.14.14.14/32`	Advertise loopback through network command inside bgp
`AGN3(config-router)#neighbor 3.3.3.3 remote-as 100`	Configure iBGP neighbor
`AGN3(config-router)#neighbor 8.8.8.8 remote-as 100`	Configure iBGP neighbor
`AGN3(config-router)#neighbor 9.9.9.9 remote-as 100`	Configure iBGP neighbor
`AGN3(config-router)#neighbor 3.3.3.3 update-source 14.14.14.14`	Update the source for the neighbor as loopback ip
`AGN3(config-router)#neighbor 8.8.8.8 update-source 14.14.14.14`	Update the source for the neighbor as loopback ip
`AGN3(config-router)#neighbor 9.9.9.9 update-source 14.14.14.14`	Update the source for the neighbor as loopback ip
`AGN3(config-router)#neighbor 3.3.3.3 fall-over bfd multihop`	Configure bfd for all the neighbors
`AGN3(config-router)#neighbor 8.8.8.8 fall-over`	Configure bfd for all the neighbors

© 2026 IP Infusion Inc. Proprietary

`bfd multihop`	
`AGN3(config-router)#neighbor 9.9.9.9 fall-over bfd multihop`	Configure bfd for all the neighbors
`AGN3(config-router)#allocate-label all`	Enable allocate label all command
`AGN3(config-router)#address-family ipv4 labeled-unicast`	Enter ipv4 labeled-unicast address family
`AGN3(config-router-af)#neighbor 3.3.3.3 activate`	Activate the neighbor under ipv4 labeled-unicast address family
`AGN3(config-router-af)#neighbor 8.8.8.8 activate`	Activate the neighbor under ipv4 labeled-unicast address family
`AGN3(config-router-af)#neighbor 9.9.9.9 activate`	Activate the neighbor under ipv4 labeled-unicast address family
`AGN3(config-router-af)#neighbor 3.3.3.3 next-hop-self`	Activate the neighbor under ipv4 labeled-unicast address family
`AGN3(config-router-af)#neighbor 8.8.8.8 next-hop-self`	Activate next hop self for the neighbor
`AGN3(config-router-af)#neighbor 9.9.9.9 next-hop-self`	Activate next hop self for the neighbor
`AGN3(config-router-af)#bgp additional-paths receive`	Enable bgp additional paths receive at global level under ipv4 labeled-unicast address-family
`AGN3(config-router-af)#neighbor 9.9.9.9 additional-paths disable`	Disable additional paths for neighbor 9.9.9.9
`AGN3(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`AGN3(config-router)#exit`	Exit Router BGP mode.
`AGN3(config)#commit`	Apply commit
`AGN3(config)#exit`	Exit Configure mode.

AGN1 Configurations

`AGN1(config)#interface lo`	Enter interface lo mode
`AGN1(config-if)#ip address 1.1.1.1/32`	Configure ip address for loopback interface
`AGN1(config-if)#exit`	Exit interface loopback mode
`AGN1(config)#router ldp`	Configure router ldp
`AGN1(config-router)#transport-address ipv4 1.1.1.1`	Configure transport address as loopback
`AGN1(config-router)#exit`	Exit the router ldp mode
`AGN1(config-if)#interface eth1`	Enter interface mode
`AGN1(config-if)#ip address 10.1.1.156/24`	Configure ip address on the interface
`AGN1(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`AGN1(config-if)#label-switching`	Enable label-switching on the interface
`AGN1(config-if)#exit`	Exit the interface mode

© 2026 IP Infusion Inc. Proprietary

`AGN1(config)#router ospf 1`	Configure router ospf instance 1
`AGN1(config-router)#network 1.1.1.1/32 area 0`	Add the loopback network into ospf instance 1
`AGN1(config-router)#network 10.1.1.0/24 area 0`	Add the ip address into ospf instance 1
`AGN1(config-router)#exit`	Exit router ospf mode
`AGN1(config)#router bgp 100`	Enter router bgp mode
`AGN1(config-router)#network 1.1.1.1/32`	Advertise loopback through network command inside bgp
`AGN1(config-router)#neighbor 3.3.3.3 remote-as 100`	Configure iBGP neighbor
`AGN1(config-router)#neighbor 8.8.8.8 remote-as 100`	Configure iBGP neighbor
`AGN1(config-router)#neighbor 9.9.9.9 remote-as 100`	Configure iBGP neighbor
`AGN1(config-router)#neighbor 3.3.3.3 update-source 1.1.1.1`	Update the source for the neighbor as loopback ip
`AGN1(config-router)#neighbor 8.8.8.8 update-source 1.1.1.1`	Update the source for the neighbor as loopback ip
`AGN1(config-router)#neighbor 9.9.9.9 update-source 1.1.1.1`	Update the source for the neighbor as loopback ip
`AGN1(config-router)#neighbor 3.3.3.3 fall-over bfd multihop`	Configure bfd for all the neighbors
`AGN1(config-router)#neighbor 8.8.8.8 fall-over bfd multihop`	Configure bfd for all the neighbors
`AGN1(config-router)#neighbor 9.9.9.9 fall-over bfd multihop`	Configure bfd for all the neighbors
`AGN1(config-router)#allocate-label all`	Enable allocate label all command
`AGN1(config-router)#address-family ipv4 labeled-unicast`	Enter ipv4 labeled-unicast address family
`AGN1(config-router-af)#neighbor 3.3.3.3 activate`	Activate the neighbor under ipv4 labeled-unicast address family
`AGN1(config-router-af)#neighbor 8.8.8.8 activate`	Activate the neighbor under ipv4 labeled-unicast address family
`AGN1(config-router-af)#neighbor 9.9.9.9 activate`	Activate the neighbor under ipv4 labeled-unicast address family
`AGN1(config-router-af)#neighbor 3.3.3.3 next-hop-self`	Activate next hop self for the neighbor
`AGN1(config-router-af)#neighbor 8.8.8.8 next-hop-self`	Activate next hop self for the neighbor
`AGN1(config-router-af)#neighbor 9.9.9.9 next-hop-self`	Activate next hop self for the neighbor
`AGN1(config-router-af)#neighbor 3.3.3.3 additional-paths receive`	Enable bgp additional paths receive at neighbor level under ipv4 labeled-unicast address-family
`AGN1(config-router-af)#neighbor 8.8.8.8 additional-paths receive`	Enable bgp additional paths receive at neighbor level under ipv4 labeled-unicast address-family

© 2026 IP Infusion Inc. Proprietary

`AGN1(config-router-af)#neighbor 9.9.9.9 additional-paths receive`	Enable bgp additional paths receive at neighbor level under ipv4 labeled-unicast address-family
`AGN1(config-router-af)#exit-address-family`	Exit bgp address-family mode.
`AGN1(config-router)#exit`	Exit Router BGP mode.
`AGN1(config)#commit`	Apply commit
`AGN1(config)#exit`	Exit Configure mode.

P1 Configurations

`P1(config)#interface lo`	Enter interface lo mode
`P1(config-if)#ip address 2.2.2.2/32`	Configure ip address for loopback interface
`P1(config-if)#exit`	Exit interface loopback mode
`P1(config)#router ldp`	Configure ldp on the router
`P1(config-router)#transport-address ipv4 2.2.2.2`	Enable loopback ip as transport address
`P1(config-router)#exit`	Exit the router ldp mode
`P1(config)#interface eth1`	Enter interface mode
`P1(config-if)#ip address 10.1.1.157/24`	Configure ip address on the interface
`P1(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`P1(config-if)#label-switching`	Enable label switching on the interface
`P1(config-if)#exit`	Exit the interface mode
`P1(config)#interface eth2`	Enter interface mode
`P1(config-if)#ip address 20.1.1.157/24`	Configure ip address on the interface
`P1(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`P1(config-if)#label-switching`	Enable label switching on the interface
`P1(config-if)#exit`	Exit the interface mode
`P1(config)#interface eth3`	Enter interface mode
`P1(config-if)#ip address 23.1.1.157/24`	Configure ip address on the interface
`P1(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`P1(config-if)#label-switching`	Enable label switching on the interface
`P1(config-if)#exit`	Exit the interface mode
`P1(config)#interface eth4`	Enter interface mode
`P1(config-if)#ip address 21.1.1.157/24`	Configure ip address on the interface
`P1(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`P1(config-if)#label-switching`	Enable label switching on the interface
`P1(config-if)#exit`	Exit the interface mode
`P1(config)#interface eth5`	Enter interface mode
`P1(config-if)#ip address 41.1.1.157/24`	Configure ip address on the interface

© 2026 IP Infusion Inc. Proprietary

P1(config-if)#enable-ldp ipv4	Enable ldp on the interface
P1(config-if)#label-switching	Enable label switching on the interface
P1(config-if)#exit	Exit the interface mode
P1(config)#router ospf 1	Enter router ospf mode with instance as 1
P1(config-router)#network 10.1.1.0/24 area 0.0.0.0	Add interface into ospf instance 1 with area 0 using network command
P1(config-router)#network 20.1.1.0/24 area 0.0.0.0	Add interface into ospf instance 1 with area 0 using network command
P1(config-router)#network 21.1.1.0/24 area 0.0.0.0	Add interface into ospf instance 1 with area 0 using network command
P1(config-router)#network 41.1.1.0/24 area 0.0.0.0	Add interface into ospf instance 1 with area 0 using network command
P1(config-router)#network 23.1.1.0/24 area 0.0.0.0	Add interface into ospf instance 1 with area 0 using network command
P1(config-router)#network 2.2.2.2/32 area 0.0.0.0	Add interface loopback into ospf instance 1 with area 0 using network command
P1(config-router)#exit	Exit Router ospf mode.
P1(config)#commit	Apply commit
P1(config)#exit	Exit Configure mode.

BR3 Configurations

BR3(config)#enable ext-ospf-multi-inst	Enable ospf multi instance
BR3(config)#interface lo	Enter interface lo mode
BR3(config-if)#ip address 8.8.8.8/32	Configure ip address for loopback interface
BR3(config-if)#exit	Exit the interface mode
BR3(config)#router ldp	Configure ldp on the router
BR3(config-router)#transport-address ipv4 8.8.8.80	Enable loopback ip as transport address
BR3(config-router)#exit	Exit the router ldp mode
BR3(config)#interface eth3	Enter interface mode
BR3(config-if)#ip address 22.1.1.165/24	Configure ip address on the interface
BR3(config-if)#enable-ldp ipv4	Enable ldp on the interface
BR3(config-if)#label-switching	Enable label switching on the interface
BR3(config-if)#exit	Exit the interface mode
BR3(config)#interface eth4	Enter interface mode
BR3(config-if)#ip address 21.1.1.165/24	Configure ip address on the interface
BR3(config-if)#enable-ldp ipv4	Enable ldp on the interface
BR3(config-if)#label-switching	Enable label switching on the interface

© 2026 IP Infusion Inc. Proprietary

BR3(config-if)#exit	Exit the interface mode
BR3(config)#router ospf 1	Configure router ospf on BR3 with instance id 1
BR3(config-router)#network 8.8.8.8/32 area 0	Add the loopback network into instance 1
BR3(config-router)#network 21.1.1.165/24 area 0	Add the network as part of ospf instance id 1
BR3(config-router)#exit	Exit the router ospf mode
BR3(config)#router ospf 2	Configure router ospf which will be part of core area with instance id as 2
BR3(config-router)#network 8.8.8.8/32 area 0.0.0.0 instance-id 2	Add interface loopback into ospf instance 2 with area 0 using network command
BR3(config-router)#network 22.1.1.165/24 area 0	Add the network as part of ospf instance id 2
BR3(config-router)#exit	Exit the router ospf mode
BR3(config)#router bgp 100	Configure router bgp
BR3(config-router)#neighbor 1.1.1.1 remote-as 100	Configure ibgp neighbor
BR3(config-router)#neighbor 14.14.14.14 remote-as 100	Configure ibgp neighbor
BR3(config-router)#neighbor 5.5.5.5 remote-as 100	Configure ibgp neighbor
BR3(config-router)#neighbor 10.10.10.10 remote-as 100	Configure ibgp neighbor
BR3(config-router)#neighbor 11.11.11.11 remote-as 100	Configure ibgp neighbor
BR3(config-router)#neighbor 12.12.12.12 remote-as 100	Configure ibgp neighbor
BR3(config-router)#neighbor 1.1.1.1 update-source 8.8.8.8	Configure update source for the neighbor with loopback address
BR3(config-router)#neighbor 14.14.14.14 update-source 8.8.8.8	Configure update source for the neighbor with loopback address
BR3(config-router)#neighbor 5.5.5.5 update-source 8.8.8.8	Configure update source for the neighbor with loopback address
BR3(config-router)#neighbor 10.10.10.10 update-source 8.8.8.8	Configure update source for the neighbor with loopback address
BR3(config-router)#neighbor 11.11.11.11 update-source 8.8.8.8	Configure update source for the neighbor with loopback address
BR3(config-router)#neighbor 12.12.12.12 update-source 8.8.8.8	Configure update source for the neighbor with loopback address
BR3(config-router)#neighbor 1.1.1.1 fall-over bfd multihop	Configure bfd for the neighbor
BR3(config-router)#neighbor 14.14.14.14 fall-over bfd multihop	Configure bfd for the neighbor
BR3(config-router)#neighbor 5.5.5.5 fall-over bfd multihop	Configure bfd for the neighbor
BR3(config-router)#neighbor 10.10.10.10 fall-over bfd multihop	Configure bfd for the neighbor

© 2026 IP Infusion Inc. Proprietary

BR3(config-router)#neighbor 11.11.11.11 fall-over bfd multihop	Configure bfd for the neighbor
BR3(config-router)#neighbor 12.12.12.12 fall-over bfd multihop	Configure bfd for the neighbor
BR3(config-router)#allocate-label all	Configure allocate label all command
R3(config-router)#address-family ipv4 labeled-unicast	Configure address-family ipv4 labeled unicast
BR3(config-router-af)#neighbor 1.1.1.1 activate	Activate the neighbor under address family labeled unicast
BR3(config-router-af)#neighbor 5.5.5.5 activate	Activate the neighbor under address family labeled unicast
BR3(config-router-af)#neighbor 10.10.10.10 activate	Activate the neighbor under address family labeled unicast
BR3(config-router-af)#neighbor 11.11.11.11 activate	Activate the neighbor under address family labeled unicast
BR3(config-router-af)#neighbor 12.12.12.12 activate	Activate the neighbor under address family labeled unicast
BR3(config-router-af)#neighbor 14.14.14.14 activate	Activate the neighbor under address family labeled unicast
BR3(config-router-af)#neighbor 1.1.1.1 route-reflector-client	Configure neighbor as route reflector client
BR3(config-router-af)#neighbor 14.14.14.14 route-reflector-client	Configure neighbor as route reflector client
BR3(config-router-af)#neighbor 5.5.5.5 next-hop-self	Configure next-hop self for the neighbor
BR3(config-router-af)#neighbor 10.10.10.10 next-hop-self	Configure next-hop self for the neighbor
BR3(config-router-af)#neighbor 11.11.11.11 next-hop-self	Configure next-hop self for the neighbor
BR3(config-router-af)#neighbor 12.12.12.12 next-hop-self	Configure next-hop self for the neighbor
BR3(config-router-af)#neighbor 1.1.1.1 additional-paths send	Enable neighbor level additional paths send for per neighbor
BR3(config-router-af)#neighbor 1.1.1.1 advertise additional-paths best 2	Enable selection of best 2 paths to be advertised to neighbor using neighbor level commands
BR3(config-router-af)#exit-address-family	Exit BGP address-family ipv4-labeled unicast mode
BR3(config-router)#end	Return to configure terminal mode

BR1 Configurations

BR1(config)#enable ext-ospf-multi-inst	Enable ospf multi instance
BR1(config)#interface lo	Enter interface lo mode

`BR1(config-if)#ip address 3.3.3.3/32`	Configure ip address for loopback interface
`BR1(config-if)#exit`	Exit the interface mode
`BR1(config)#router ldp`	Configure ldp on the router
`BR1(config-router)#transport-address ipv4 3.3.3.3 0`	Enable loopback ip as transport address
`BR1(config-router)#exit`	Exit the router ldp mode
`BR1(config)#interface eth1`	Enter interface mode
`BR1(config-if)#ip address 30.1.1.158/24`	Configure ip address on the interface
`BR1(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`BR1(config-if)#label-switching`	Enable label switching on the interface
`BR1(config-if)#exit`	Exit the interface mode
`BR1(config)#interface eth2`	Enter interface mode
`BR1(config-if)#ip address 20.1.1.158/24`	Configure ip address on the interface
`BR1(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`BR1(config-if)#label-switching`	Enable label switching on the interface
`BR1(config-if)#exit`	Exit the interface mode
`BR1(config)#router ospf 1`	Configure router ospf with instance id 1
`BR1(config-router)#network 3.3.3.3/32 area 0`	Add the loopback network into instance 1
`BR1(config-router)#network 20.1.1.0/24 area 0`	Add the network as part of ospf instance id 1
`BR1(config-router)#exit`	Exit the router ospf mode
`BR1(config)#router ospf 2`	Configure router ospf which will be part of core area with instance id as 2
`BR1(config-router)#network 3.3.3.3/32 area 0.0.0.0 instance-id 2`	Add interface loopback into ospf instance 2 with area 0 using network command
`BR1(config-router)#network 30.1.1.0/24 area 0`	Add the network as part of ospf instance id 2
`BR1(config-router)#exit`	Exit the router ospf mode
`BR1(config)#route-map BEST2`	Create a route-map to select best 2 paths
`BR1(config-route-map)#match additional-paths advertise-set select best 2`	Select best 2 paths to advertise to the neighbor
`BR1(config-route-map)#exit`	Exit the route-map mode
`BR1(config)#router bgp 100`	Configure router bgp
`BR1(config-router)#neighbor 1.1.1.1 remote-as 100`	Configure ibgp neighbor
`BR1(config-router)#neighbor 14.14.14.14 remote-as 100`	Configure ibgp neighbor
`BR1(config-router)#neighbor 5.5.5.5 remote-as 100`	Configure ibgp neighbor
`BR1(config-router)#neighbor 10.10.10.10 remote-as 100`	Configure ibgp neighbor
`BR1(config-router)#neighbor 11.11.11.11 remote-as 100`	Configure ibgp neighbor

© 2026 IP Infusion Inc. Proprietary

BR1(config-router)#neighbor 12.12.12.12 remote-as 100	Configure ibgp neighbor
BR1(config-router)#neighbor 1.1.1.1 update-source 3.3.3.3	Configure update source for the neighbor with loopback address
BR1(config-router)#neighbor 14.14.14.14 update-source 3.3.3.3	Configure update source for the neighbor with loopback address
BR1(config-router)#neighbor 5.5.5.5 update-source 3.3.3.3	Configure update source for the neighbor with loopback address
BR1(config-router)#neighbor 10.10.10.10 update-source 3.3.3.3	Configure update source for the neighbor with loopback address
BR1(config-router)#neighbor 11.11.11.11 update-source 3.3.3.3	Configure update source for the neighbor with loopback address
BR1(config-router)#neighbor 12.12.12.12 update-source 3.3.3.3	Configure update source for the neighbor with loopback address
BR1(config-router)#neighbor 1.1.1.1 fall-over bfd multihop	Configure bfd for the neighbor
BR1(config-router)#neighbor 14.14.14.14 fall-over bfd multihop	Configure bfd for the neighbor
BR1(config-router)#neighbor 5.5.5.5 fall-over bfd multihop	Configure bfd for the neighbor
BR1(config-router)#neighbor 10.10.10.10 fall-over bfd multihop	Configure bfd for the neighbor
BR1(config-router)#neighbor 11.11.11.11 fall-over bfd multihop	Configure bfd for the neighbor
BR1(config-router)#neighbor 12.12.12.12 fall-over bfd multihop	Configure bfd for the neighbor
BR1(config-router)#allocate-label all	Configure allocate label all command
BR1(config-router)#address-family ipv4 labeled-unicast	Configure address-family ipv4 labeled unicast
BR1(config-router-af)#neighbor 1.1.1.1 activate	Activate the neighbor under address family labeled unicast
BR1(config-router-af)#neighbor 5.5.5.5 activate	Activate the neighbor under address family labeled unicast
BR1(config-router-af)#neighbor 10.10.10.10 activate	Activate the neighbor under address family labeled unicast
BR1(config-router-af)#neighbor 11.11.11.11 activate	Activate the neighbor under address family labeled unicast
BR1(config-router-af)#neighbor 12.12.12.12 activate	Activate the neighbor under address family labeled unicast
BR1(config-router-af)#neighbor 14.14.14.14 activate	Activate the neighbor under address family labeled unicast
BR1(config-router-af)#neighbor 1.1.1.1 route-reflector-client	Configure neighbor as route reflector client
BR1(config-router-af)#neighbor 14.14.14.14 route-	Configure neighbor as route reflector client

reflector-client	
BR1(config-router-af)#neighbor 5.5.5.5 next-hop-self	Configure next-hop self for the neighbor
BR1(config-router-af)#neighbor 10.10.10.10 next-hop-self	Configure next-hop self for the neighbor
BR1(config-router-af)#neighbor 11.11.11.11 next-hop-self	Configure next-hop self for the neighbor
BR1(config-router-af)#neighbor 12.12.12.12 next-hop-self	Configure next-hop self for the neighbor
BR1(config-router-af)#neighbor 1.1.1.1 additional-paths send	Enable additional paths send capability on neighbor level under ipv4 labeled unicast address family
BR1(config-router-af)#neighbor 14.14.14.14 additional-paths send	Enable additional paths send capability on neighbor level under ipv4 labeled unicast address family
BR1(config-router-af)#neighbor 1.1.1.1 advertise additional-paths all	Enable selection of all paths to be advertised to neighbor using neighbor level commands
BR1(config-router-af)#neighbor 14.14.14.14 route-map BEST2 out	Enable selection of all paths to be advertised to neighbor using neighbor level commands
BR1(config-router-af)#exit-address-family	Exit BGP address-family ipv4-labeled unicast mode
BR1(config-router)#end	Return to configure terminal mode

BR4 Configurations

BR4(config)#enable ext-ospf-multi-inst	Enable ospf multi instance
BR4(config)#interface lo	Enter interface lo mode
BR4(config-if)#ip address 9.9.9.9/32	Configure ip address for loopback interface
BR4(config-if)#exit	Exit the interface mode
BR4(config)#access-list 1 permit 10.10.10.10 0.0.0.0	Configure the access list to allow the BR5 loopback ip
BR4(config)#access-list 2 permit 11.11.11.11 0.0.0.0	Configure the access list to allow the BR6 loopback ip
BR4(config)#access-list 3 permit 5.5.5.5 0.0.0.0	Configure the access list to allow the BR2 loopback ip
BR4(config)#access-list 4 permit 12.12.12.12 0.0.0.0	Configure the access list to allow the BR7 loopback ip
BR4(config)#route-map A permit 10	Configure route-map A and permit
BR4(config-route-map)#match ip address 1	Match ip address 1
BR4(config-route-map)#exit	Exit the route-map
BR4(config)#route-map A permit 20	Configure route-map A and permit

© 2026 IP Infusion Inc. Proprietary

`BR4(config-route-map)#match ip address 2`	Match ip address 2
`BR4(config-route-map)#exit`	Exit the route-map
`BR4(config)#route-map A permit 30`	Configure route-map A and permit
`BR4(config-route-map)#match ip address 3`	Match ip address 3
`BR4(config-route-map)#exit`	Exit the route-map
`BR4(config)#route-map A permit 40`	Configure route-map A and permit
`BR4(config-route-map)#match ip address 4`	Match ip address 4
`BR4(config-route-map)#exit`	Exit the route-map
`BR4(config)#router ldp`	Configure ldp on the router
`BR4(config-router)#transport-address ipv4 9.9.9.9 0`	Enable loopback ip as transport address
`BR4(config-router)#exit`	Exit the router ldp mode
`BR4(config)#interface eth3`	Enter interface mode
`BR4(config-if)#ip address 23.1.1.166/24`	Configure ip address on the interface
`BR4(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`BR4(config-if)#label-switching`	Enable label switching on the interface
`BR4(config-if)#exit`	Exit the interface mode
`BR4(config)#interface eth4`	Enter interface mode
`BR4(config-if)#ip address 24.1.1.166/24`	Configure ip address on the interface
`BR4(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`BR4(config-if)#label-switching`	Enable label switching on the interface
`BR4(config-if)#exit`	Exit the interface mode
`BR4(config)#router ospf 1`	Configure router ospf with instance id 1
`BR4(config-router)#network 9.9.9.9/32 area 0`	Add the loopback network into instance 1
`BR4(config-router)#network 23.1.1.0/24 area 0`	Add the network as part of ospf instance id 1
`BR4(config-router)#redistribute ospf 2 route-map A`	Redistribute the selected routes from ospf instance 2 into instance 1
`BR4(config-router)#exit`	Exit the router ospf mode
`BR4(config)#router ospf 2`	Configure router ospf which will be part of core area with instance id as 2
`BR4(config-router)#network 9.9.9.9/32 area 0.0.0.0 instance-id 2`	Add interface loopback into ospf instance 2 with area 0 using network command
`BR4(config-router)#network 24.1.1.0/24 area 0`	Add the network as part of ospf instance id 2
`BR4(config-router)#exit`	Exit the router ospf mode
`BR4(config)#router bgp 100`	Configure router bgp
`BR4(config-router)#neighbor 1.1.1.1 remote-as 100`	Configure ibgp neighbor
`BR4(config-router)#neighbor 14.14.14.14 remote-as 100`	Configure ibgp neighbor

© 2026 IP Infusion Inc. Proprietary

BR4(config-router)#neighbor 5.5.5.5 remote-as 100	Configure ibgp neighbor
BR4(config-router)#neighbor 10.10.10.10 remote-as 100	Configure ibgp neighbor
BR4(config-router)#neighbor 11.11.11.11 remote-as 100	Configure ibgp neighbor
BR4(config-router)#neighbor 12.12.12.12 remote-as 100	Configure ibgp neighbor
BR4(config-router)#neighbor 1.1.1.1 update-source 9.9.9.9	Configure update source for the neighbor with loopback address
BR4(config-router)#neighbor 14.14.14.14 update-source 9.9.9.9	Configure update source for the neighbor with loopback address
BR4(config-router)#neighbor 5.5.5.5 update-source 9.9.9.9	Configure update source for the neighbor with loopback address
BR4(config-router)#neighbor 10.10.10.10 update-source 3.3.3.3	Configure update source for the neighbor with loopback address
BR4(config-router)#neighbor 11.11.11.11 update-source 9.9.9.9	Configure update source for the neighbor with loopback address
BR4(config-router)#neighbor 12.12.12.12 update-source 9.9.9.9	Configure update source for the neighbor with loopback address
BR4(config-router)#neighbor 1.1.1.1 fall-over bfd multihop	Configure bfd for the neighbor
BR4(config-router)#neighbor 14.14.14.14 fall-over bfd multihop	Configure bfd for the neighbor
BR4(config-router)#neighbor 5.5.5.5 fall-over bfd multihop	Configure bfd for the neighbor
BR4(config-router)#neighbor 10.10.10.10 fall-over bfd multihop	Configure bfd for the neighbor
BR4(config-router)#neighbor 11.11.11.11 fall-over bfd multihop	Configure bfd for the neighbor
BR4(config-router)#neighbor 12.12.12.12 fall-over bfd multihop	Configure bfd for the neighbor
BR4(config-router)#allocate-label all	Configure allocate label all command
BR4(config-router)#address-family ipv4 labeled-unicast	Configure address-family ipv4 labeled unicast
BR4(config-router-af)#neighbor 1.1.1.1 activate	Activate the neighbor under address family labeled unicast
BR4(config-router-af)#neighbor 5.5.5.5 activate	Activate the neighbor under address family labeled unicast
BR4(config-router-af)#neighbor 10.10.10.10 activate	Activate the neighbor under address family labeled unicast
BR4(config-router-af)#neighbor 11.11.11.11 activate	Activate the neighbor under address family labeled unicast
BR4(config-router-af)#neighbor 12.12.12.12 activate	Activate the neighbor under address family labeled unicast

© 2026 IP Infusion Inc. Proprietary

BR4(config-router-af)#neighbor 14.14.14.14 activate	Activate the neighbor under address family labeled unicast
BR4(config-router-af)#neighbor 1.1.1.1 route-reflector-client	Configure neighbor as route reflector client
BR4(config-router-af)#neighbor 14.14.14.14 route-reflector-client	Configure neighbor as route reflector client
BR4(config-router-af)#neighbor 5.5.5.5 next-hop-self	Configure next-hop self for the neighbor
BR4(config-router-af)#neighbor 10.10.10.10 next-hop-self	Configure next-hop self for the neighbor
BR4(config-router-af)#neighbor 11.11.11.11 next-hop-self	Configure next-hop self for the neighbor
BR4(config-router-af)#neighbor 12.12.12.12 next-hop-self	Configure next-hop self for the neighbor
BR4(config-router-af)#bgp additional-paths send	Enable additional paths send capability at global level under address-family labeled-unicast
BR4(config-router-af)#neighbor 1.1.1.1 advertise additional-paths best 3	Enable selection of best 3 paths per neighbor using neighbor level commands
BR4(config-router-af)#neighbor 14.14.14.14 advertise additional-paths all	Enable selection of all paths to be advertised to neighbor using neighbor level commands
BR4(config-router-af)#exit-address-family	Exit BGP address-family ipv4-labeled unicast mode
BR4(config-router)#end	Return to configure terminal mode

P2 Configurations

P2(config)#interface lo	Enter interface lo mode
P2(config-if)#ip address 4.4.4.4/32	Configure ip address for loopback interface
P2(config-if)#exit	Exit interface loopback mode
P2(config)#router ldp	Configure ldp on the router
P2(config-router)#transport-address ipv4 4.4.4.40	Enable loopback ip as transport address
P2(config-router)#exit	Exit the router ldp mode
P2(config)#interface eth1	Enter interface mode
P2(config-if)#ip address 30.1.1.159/24	Configure ip address on the interface
P2(config-if)#enable-ldp ipv4	Enable ldp on the interface
P2(config-if)#label-switching	Enable label switching on the interface
P2(config-if)#exit	Exit the interface mode
P2(config)#interface eth2	Enter interface mode
P2(config-if)#ip address 40.1.1.159/24	Configure ip address on the interface
P2(config-if)#enable-ldp ipv4	Enable ldp on the interface

© 2026 IP Infusion Inc. Proprietary

`P2(config-if)#label-switching`	Enable label switching on the interface
`P2(config-if)#exit`	Exit the interface mode
`P2(config)#interface eth3`	Enter interface mode
`P2(config-if)#ip address 22.1.1.159/24`	Configure ip address on the interface
`P2(config-if)#enable-ldp ipv4`	Enable rsvp on the interface
`P2(config-if)#label-switching`	Enable label switching on the interface
`P2(config-if)#exit`	Exit the interface mode
`P2(config)#interface eth4`	Enter interface mode
`P2(config-if)#ip address 24.1.1.159/24`	Configure ip address on the interface
`P2(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`P2(config-if)#label-switching`	Enable label switching on the interface
`P2(config-if)#exit`	Exit the interface mode
`P2(config)#interface eth5`	Enter interface mode
`P2(config-if)#ip address 25.1.1.159/24`	Configure ip address on the interface
`P2(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`P2(config-if)#label-switching`	Enable label switching on the interface
`P2(config-if)#exit`	Exit the interface mode
`P2(config)#interface eth6`	Enter interface mode
`P2(config-if)#ip address 27.1.1.159/24`	Configure ip address on the interface
`P2(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`P2(config-if)#label-switching`	Enable label switching on the interface
`P2(config-if)#exit`	Exit the interface mode
`P2(config)#interface eth7`	Enter interface mode
`P2(config-if)#ip address 29.1.1.159/24`	Configure ip address on the interface
`P2(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`P2(config-if)#label-switching`	Enable label switching on the interface
`P2(config-if)#exit`	Exit the interface mode
`P2(config)#router ospf 2`	Configure Router ospf with instance id 2
`P2(config-router)#network 30.1.1.0/24 area 0.0.0.0`	Add interface into ospf instance 2 with area 0 using network command
`P2(config-router)#network 40.1.1.0/24 area 0.0.0.0`	Add interface into ospf instance 2 with area 0 using network command
`P2(config-router)#network 22.1.1.0/24 area 0.0.0.0`	Add interface into ospf instance 2 with area 0 using network command
`P2(config-router)#network 24.1.1.0/24 area 0.0.0.0`	Add interface into ospf instance 2 with area 0 using network command
`P2(config-router)#network 25.1.1.0/24 area 0.0.0.0`	Add interface into ospf instance 2 with area 0

	using network command
P2(config-router)#network 27.1.1.0/24 area 0.0.0.0	Add interface into ospf instance 2 with area 0 using network command
P2(config-router)#network 29.1.1.0/24 area 0.0.0.0	Add interface into ospf instance 2 with area 0 using network command
P2(config-router)#network 4.4.4.4/32 area 0.0.0.0	Add interface loopback into ospf instance 2 with area 0 using network command
P2(config-router)#exit	Exit router ospf mode

BR5 Configurations

BR5(config)#enable ext-ospf-multi-inst	Enable ospf multi instance
BR5(config)#interface lo	Enter interface lo mode
BR5(config-if)#ip address 10.10.10.10/32	Configure ip address for loopback interface
BR5(config-if)#exit	Exit the interface mode
BR5(config)#router ldp	Configure ldp on the router
BR5(config-router)#transport-address ipv4 10.10.10.10 0	Enable loopback ip as transport address
BR5(config-router)#exit	Exit router ldp mode
BR5(config)#interface eth3	Enter interface mode
BR5(config-if)#ip address 26.1.1.167/24	Configure ip address on the interface
BR5(config-if)#enable-ldp ipv4	Enable ldp on the interface
BR5(config-if)#label-switching	Enable label switching on the interface
BR5(config-if)#exit	Exit the interface mode
BR5(config)#interface eth5	Enter interface mode
BR5(config-if)#ip address 25.1.1.167/24	Configure ip address on the interface
BR5(config-if)#enable-ldp ipv4	Enable ldp on the interface
BR5(config-if)#label-switching	Enable label switching on the interface
BR5(config-if)#exit	Exit the interface mode
BR5(config)#router ospf 2	Configure router ospf with instance id 2
BR5(config-router)#network 10.10.10.10/32 area 0	Add the loopback network into instance 2
BR5(config-router)#network 25.1.1.0/24 area 0	Add the network as part of ospf instance id 2
BR5(config-router)#exit	Exit the router ospf mode
BR5(config)#router ospf 3	Configure router ospf which will be part of aggregation area 2 with instance id as 3
BR5(config-router)#network 10.10.10.10/32 area 0.0.0.0 instance-id 3	Add interface loopback into ospf instance 3 with area 0 using network command
BR5(config-router)#network 26.1.1.0/24 area 0	Add the network as part of ospf instance id 3

© 2026 IP Infusion Inc. Proprietary

BR5(config-router)#exit	Exit the router ospf mode
BR5(config)#router bgp 100	Configure router bgp
BR5(config-router)#neighbor 3.3.3.3 remote-as 100	Configure ibgp neighbor
BR5(config-router)#neighbor 7.7.7.7 remote-as 100	Configure ibgp neighbor
BR5(config-router)#neighbor 8.8.8.8 remote-as 100	Configure ibgp neighbor
BR5(config-router)#neighbor 9.9.9.9 remote-as 100	Configure ibgp neighbor
BR5(config-router)#neighbor 13.13.13.13 remote-as 100	Configure ibgp neighbor
BR5(config-router)#neighbor 15.15.15.15 remote-as 100	Configure ibgp neighbor
BR5(config-router)#neighbor 3.3.3.3 update-source 10.10.10.10	Configure update source for the neighbor with loopback address
BR5(config-router)#neighbor 7.7.7.7 update-source 10.10.10.10	Configure update source for the neighbor with loopback address
BR5(config-router)#neighbor 8.8.8.8 update-source 10.10.10.10	Configure update source for the neighbor with loopback address
BR5(config-router)#neighbor 9.9.9.9 update-source 10.10.10.10	Configure update source for the neighbor with loopback address
BR5(config-router)#neighbor 13.13.13.13 update-source 10.10.10.10	Configure update source for the neighbor with loopback address
BR5(config-router)#neighbor 15.15.15.15 update-source 10.10.10.10	Configure update source for the neighbor with loopback address
BR5(config-router)#neighbor 3.3.3.3 fall-over bfd multihop	Configure bfd for the neighbor
BR5(config-router)#neighbor 7.7.7.7 fall-over bfd multihop	Configure bfd for the neighbor
BR5(config-router)#neighbor 8.8.8.8 fall-over bfd multihop	Configure bfd for the neighbor
BR5(config-router)#neighbor 9.9.9.9 fall-over bfd multihop	Configure bfd for the neighbor
BR5(config-router)#neighbor 13.13.13.13 fall-over bfd multihop	Configure bfd for the neighbor
BR5(config-router)#neighbor 15.15.15.15 fall-over bfd multihop	Configure bfd for the neighbor
BR5(config-router)#allocate-label all	Configure allocate label all command
BR5(config-router)#address-family ipv4 labeled-unicast	Configure address-family ipv4 labeled unicast
BR5(config-router-af)#neighbor 3.3.3.3 activate	Activate the neighbor under address family labeled unicast
BR5(config-router-af)#neighbor 7.7.7.7 activate	Activate the neighbor under address family labeled unicast
BR5(config-router-af)#neighbor 8.8.8.8 activate	Activate the neighbor under address family labeled unicast

© 2026 IP Infusion Inc. Proprietary

BR5(config-router-af)#neighbor 9.9.9.9 activate	Activate the neighbor under address family labeled unicast
BR5(config-router-af)#neighbor 13.13.13.13 activate	Activate the neighbor under address family labeled unicast
BR5(config-router-af)#neighbor 15.15.15.15 activate	Activate the neighbor under address family labeled unicast
BR5(config-router-af)#neighbor 13.13.13.13 route-reflector-client	Configure neighbor as route reflector client
BR5(config-router-af)#neighbor 7.7.7.7 route-reflector-client	Configure neighbor as route reflector client
BR5(config-router-af)#neighbor 15.15.15.15 next-hop-self	Configure next-hop self for the neighbor
BR5(config-router-af)#neighbor 3.3.3.3 next-hop-self	Configure next-hop self for the neighbor
BR5(config-router-af)#neighbor 8.8.8.8 next-hop-self	Configure next-hop self for the neighbor
BR5(config-router-af)#neighbor 9.9.9.9 next-hop-self	Configure next-hop self for the neighbor
BR5(config-router-af)#bgp additional-paths send	Enable additional paths send capability on neighbor level under ipv4 labeled unicast address family
BR5(config-router-af)#bgp additional-paths select all	Select all the available paths to advertise to the next router
BR5(config-router)#exit	Exit BGP router mode

BR6 Configurations

BR6(config)#enable ext-ospf-multi-inst	Enable ospf multi instance
BR6(config)#interface lo	Enter interface lo mode
BR6(config-if)#ip address 11.11.11.11/32	Configure ip address for loopback interface
BR6(config-if)#exit	Exit the interface mode
BR6(config)#router ldp	Configure ldp on the router
BR6(config-router)#transport-address ipv4 11.11.11.11	Enable loopback ip as transport address
BR5(config-router)#exit	Exit router ldp mode
BR6(config)#interface eth4	Enter interface mode
BR6(config-if)#ip address 28.1.1.168/24	Configure ip address on the interface
BR6(config-if)#enable-ldp ipv4	Enable ldp on the interface
BR6(config-if)#label-switching	Enable label switching on the interface
BR6(config-if)#exit	Exit the interface mode
BR6(config)#interface eth6	Enter interface mode
BR6(config-if)#ip address 27.1.1.168/24	Configure ip address on the interface

© 2026 IP Infusion Inc. Proprietary

`BR6(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`BR6(config-if)#label-switching`	Enable label switching on the interface
`BR6(config-if)#exit`	Exit the interface mode
`BR6(config)#router ospf 2`	Configure router ospf with instance id 2
`BR6(config-router)#network 11.11.11.11/32 area 0`	Add the loopback network into instance 2
`BR6(config-router)#network 27.1.1.0/24 area 0`	Add the network as part of ospf instance id 2
`BR6(config-router)#exit`	Exit the router ospf mode
`BR6(config)#router ospf 3`	Configure router ospf which will be part of aggregation 2 with instance id as 3
`BR6(config-router)#network 11.11.11.11/32 area 0.0.0.0 instance-id 3`	Add interface loopback into ospf instance 3 with area 0 using network command
`BR6(config-router)#network 28.1.1.0/24 area 0`	Add the network as part of ospf instance id 3
`BR6(config-router)#exit`	Exit the router ospf mode
`BR6(config)#router bgp 100`	Configure router bgp
`BR6(config-router)#neighbor 3.3.3.3 remote-as 100`	Configure ibgp neighbor
`BR6(config-router)#neighbor 7.7.7.7 remote-as 100`	Configure ibgp neighbor
`BR6(config-router)#neighbor 8.8.8.8 remote-as 100`	Configure ibgp neighbor
`BR6(config-router)#neighbor 9.9.9.9 remote-as 100`	Configure ibgp neighbor
`BR6(config-router)#neighbor 13.13.13.13 remote-as 100`	Configure ibgp neighbor
`BR6(config-router)#neighbor 15.15.15.15 remote-as 100`	Configure ibgp neighbor
`BR6(config-router)#neighbor 3.3.3.3 update-source 11.11.11.11`	Configure update source for the neighbor with loopback address
`BR6(config-router)#neighbor 7.7.7.7 update-source 11.11.11.11`	Configure update source for the neighbor with loopback address
`BR6(config-router)#neighbor 8.8.8.8 update-source 11.11.11.11`	Configure update source for the neighbor with loopback address
`BR6(config-router)#neighbor 9.9.9.9 update-source 11.11.11.11`	Configure update source for the neighbor with loopback address
`BR6(config-router)#neighbor 13.13.13.13 update-source 11.11.11.11`	Configure update source for the neighbor with loopback address
`BR6(config-router)#neighbor 15.15.15.15 update-source 11.11.11.11`	Configure update source for the neighbor with loopback address
`BR6(config-router)#neighbor 3.3.3.3 fall-over bfd multihop`	Configure bfd for the neighbor
`BR6(config-router)#neighbor 7.7.7.7 fall-over bfd multihop`	Configure bfd for the neighbor
`BR6(config-router)#neighbor 8.8.8.8 fall-over bfd multihop`	Configure bfd for the neighbor
`BR6(config-router)#neighbor 9.9.9.9 fall-over bfd`	Configure bfd for the neighbor

© 2026 IP Infusion Inc. Proprietary

multihop	
BR6(config-router)#neighbor 13.13.13.13 fall-over bfd multihop	Configure bfd for the neighbor
BR6(config-router)#neighbor 15.15.15.15 fall-over bfd multihop	Configure bfd for the neighbor
BR6(config-router)#allocate-label all	Configure allocate label all command
BR6(config-router)#address-family ipv4 labeled-unicast	Configure address-family ipv4 labeled unicast
BR6(config-router-af)#neighbor 3.3.3.3 activate	Activate the neighbor under address family labeled unicast
BR6(config-router-af)#neighbor 7.7.7.7 activate	Activate the neighbor under address family labeled unicast
BR6(config-router-af)#neighbor 8.8.8.8 activate	Activate the neighbor under address family labeled unicast
BR6(config-router-af)#neighbor 9.9.9.9 activate	Activate the neighbor under address family labeled unicast
BR6(config-router-af)#neighbor 13.13.13.13 activate	Activate the neighbor under address family labeled unicast
BR6(config-router-af)#neighbor 15.15.15.15 activate	Activate the neighbor under address family labeled unicast
BR6(config-router-af)#neighbor 13.13.13.13 route-reflector-client	Configure neighbor as route reflector client
BR6(config-router-af)#neighbor 7.7.7.7 route-reflector-client	Configure neighbor as route reflector client
BR6(config-router-af)#neighbor 15.15.15.15 route-reflector-client	Configure neighbor as route reflector client
BR6(config-router-af)#neighbor 15.15.15.15 next-hop-self	Configure next-hop self for the neighbor
BR6(config-router-af)#neighbor 3.3.3.3 next-hop-self	Configure next-hop self for the neighbor
BR6(config-router-af)#neighbor 8.8.8.8 next-hop-self	Configure next-hop self for the neighbor
BR6(config-router-af)#neighbor 9.9.9.9 next-hop-self	Configure next-hop self for the neighbor
BR6(config-router-af)#bgp additional-paths send	Enable additional paths send capability on neighbor level under ipv4 labeled unicast address family
BR6(config-router-af)#bgp additional-paths select best 2	Select best 2 paths from the available additional paths to advertise to the next router
BR6(config-router)#exit	Exit BGP router mode

BR2 Configurations

BR2(config)#enable ext-ospf-multi-inst	Enable ospf multi instance

`BR2(config)#interface lo`	Enter interface lo mode
`BR2(config-if)#ip address 5.5.5.5/32`	Configure ip address for loopback interface
`BR2(config-if)#exit`	Exit the interface mode
`BR2(config)#router ldp`	Configure ldp on the router
`BR2(config-router)#transport-address ipv4 5.5.5.5`	Enable loopback ip as transport address
`BR2(config-router)#exit`	Exit router ldp mode
`BR2(config)#interface eth1`	Enter interface mode
`BR2(config-if)#ip address 50.1.1.160/24`	Configure ip address on the interface
`BR2(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`BR2(config-if)#label-switching`	Enable label switching on the interface
`BR2(config-if)#exit`	Exit the interface mode
`BR2(config)#interface eth2`	Enter interface mode
`BR2(config-if)#ip address 40.1.1.160/24`	Configure ip address on the interface
`BR2(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`BR2(config-if)#label-switching`	Enable label switching on the interface
`BR2(config-if)#exit`	Exit the interface mode
`BR2(config)#router ospf 2`	Configure router ospf with instance id 2
`BR2(config-router)#network 5.5.5.5/32 area 0`	Add the loopback network into instance 2
`BR2(config-router)#network 40.1.1.0/24 area 0`	Add the network as part of ospf instance id 2
`BR2(config-router)#exit`	Exit the router ospf mode
`BR2(config)#router ospf 3`	Configure router ospf which will be part of aggregation 2 with instance id as 3
`BR2(config-router)#network 5.5.5.5/32 area 0.0.0.0 instance-id 3`	Add interface loopback into ospf instance 3 with area 0 using network command
`BR2(config-router)#network 50.1.1.0/24 area 0`	Add the network as part of ospf instance id 3
`BR2(config-router)#exit`	Exit the router ospf mode
`BR2(config)#router bgp 100`	Configure router bgp
`BR2(config-router)#neighbor 3.3.3.3 remote-as 100`	Configure ibgp neighbor
`BR2(config-router)#neighbor 7.7.7.7 remote-as 100`	Configure ibgp neighbor
`BR2(config-router)#neighbor 8.8.8.8 remote-as 100`	Configure ibgp neighbor
`BR2(config-router)#neighbor 9.9.9.9 remote-as 100`	Configure ibgp neighbor
`BR2(config-router)#neighbor 13.13.13.13 remote-as 100`	Configure ibgp neighbor
`BR2(config-router)#neighbor 15.15.15.15 remote-as 100`	Configure ibgp neighbor
`BR2(config-router)#neighbor 3.3.3.3 update-source 5.5.5.5`	Configure update source for the neighbor with loopback address
`BR2(config-router)#neighbor 7.7.7.7 update-source`	Configure update source for the neighbor with

5.5.5.5	loopback address
BR2(config-router)#neighbor 8.8.8.8 update-source 5.5.5.5	Configure update source for the neighbor with loopback address
BR2(config-router)#neighbor 9.9.9.9 update-source 5.5.5.5	Configure update source for the neighbor with loopback address
BR2(config-router)#neighbor 13.13.13.13 update-source 5.5.5.5	Configure update source for the neighbor with loopback address
BR2(config-router)#neighbor 15.15.15.15 update-source 5.5.5.5	Configure update source for the neighbor with loopback address
BR2(config-router)#neighbor 3.3.3.3 fall-over bfd multihop	Configure bfd for the neighbor
BR2(config-router)#neighbor 7.7.7.7 fall-over bfd multihop	Configure bfd for the neighbor
BR2(config-router)#neighbor 8.8.8.8 fall-over bfd multihop	Configure bfd for the neighbor
BR2(config-router)#neighbor 9.9.9.9 fall-over bfd multihop	Configure bfd for the neighbor
BR2(config-router)#neighbor 13.13.13.13 fall-over bfd multihop	Configure bfd for the neighbor
BR2(config-router)#neighbor 15.15.15.15 fall-over bfd multihop	Configure bfd for the neighbor
BR2(config-router)#allocate-label all	Configure allocate label all command
BR2(config-router)#address-family ipv4 labeled-unicast	Configure address-family ipv4 labeled unicast
BR2(config-router-af)#neighbor 3.3.3.3 activate	Activate the neighbor under address family labeled unicast
BR2(config-router-af)#neighbor 7.7.7.7 activate	Activate the neighbor under address family labeled unicast
BR2(config-router-af)#neighbor 8.8.8.8 activate	Activate the neighbor under address family labeled unicast
BR2(config-router-af)#neighbor 9.9.9.9 activate	Activate the neighbor under address family labeled unicast
BR2(config-router-af)#neighbor 13.13.13.13 activate	Activate the neighbor under address family labeled unicast
BR2(config-router-af)#neighbor 15.15.15.15 activate	Activate the neighbor under address family labeled unicast
BR2(config-router-af)#neighbor 13.13.13.13 route-reflector-client	Configure neighbor as route reflector client
BR2(config-router-af)#neighbor 7.7.7.7 route-reflector-client	Configure neighbor as route reflector client
BR2(config-router-af)#neighbor 15.15.15.15 route-reflector-client	Configure neighbor as route reflector client
BR2(config-router-af)#neighbor 15.15.15.15 next-hop-self	Configure next-hop self for the neighbor

BR2(config-router-af)#neighbor 3.3.3.3 next-hop-self	Configure next-hop self for the neighbor
BR2(config-router-af)#neighbor 8.8.8.8 next-hop-self	Configure next-hop self for the neighbor
BR2(config-router-af)#neighbor 9.9.9.9 next-hop-self	Configure next-hop self for the neighbor
BR2(config-router-af)#bgp additional-paths send	Enable additional paths send capability on neighbor level under ipv4 labeled unicast address family
BR2(config-router-af)#bgp additional-paths select best 3	Select best 3 paths from the available additional paths to advertise to the next router
BR2(config-router)#exit	Exit BGP router mode

BR7 Configurations

BR7(config)#enable ext-ospf-multi-inst	Enable ospf multi instance
BR7(config)#interface lo	Enter interface lo mode
BR7(config-if)#ip address 12.12.12.12/32	Configure ip address for loopback interface
BR7(config-if)#exit	Exit the interface mode
BR7(config)#access-list 1 permit 3.3.3.3 0.0.0.0	Configure the access list to allow the BR5 loopback ip
BR7(config)#access-list 2 permit 8.8.8.8 0.0.0.0	Configure the access list to allow the BR6 loopback ip
BR7(config)#access-list 3 permit 9.9.9.9 0.0.0.0	Configure the access list to allow the BR7 loopback ip
BR7(config)#route-map A permit 10	Configure route-map A and permit
BR7(config-route-map)#match ip address 1	Match ip address 1
BR7(config-route-map)#exit	Exit the route-map
BR7(config)#route-map A permit 20	Configure route-map A and permit
BR7(config-route-map)#match ip address 2	Match ip address 2
BR7(config-route-map)#exit	Exit the route-map
BR7(config)#route-map A permit 30	Configure route-map A and permit
BR7(config-route-map)#match ip address 3	Match ip address 3
BR7(config-route-map)#exit	Exit the route-map
BR7(config)#router ldp	Configure ldp on the router
BR7(config-router)#transport-address ipv4 12.12.12.12 0	Enable loopback ip as transport address
BR7(config-router)#exit	Exit router ldp mode
BR7(config)#interface eth5	Enter interface mode
BR7(config-if)#ip address 33.1.1.169/24	Configure ip address on the interface

© 2026 IP Infusion Inc. Proprietary

`BR7(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`BR7(config-if)#label-switching`	Enable label switching on the interface
`BR7(config-if)#exit`	Exit the interface mode
`BR7(config)#interface eth7`	Enter interface mode
`BR7(config-if)#ip address 29.1.1.169/24`	Configure ip address on the interface
`BR7(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`BR7(config-if)#label-switching`	Enable label switching on the interface
`BR7(config-if)#exit`	Exit the interface mode
`BR7(config)#router ospf 2`	Configure router ospf with instance id 2
`BR7(config-router)#network 12.12.12.12/32 area 0`	Add the loopback network into instance 2
`BR7(config-router)#network 29.1.1.0/24 area 0`	Add the network as part of ospf instance id 2
`BR7(config-router)#exit`	Exit the router ospf mode
`BR7(config)#router ospf 3`	Configure router ospf which will be part of aggregation 2 with instance id as 3
`BR7(config-router)#redistribute ospf 2 route-map A`	Redistribute selected loopbacks into ospf instance 3
`BR7(config-router)#network 12.12.12.12/32 area 0.0.0.0 instance-id 3`	Add interface loopback into ospf instance 3 with area 0 using network command
`BR7(config-router)#network 33.1.1.0/24 area 0`	Add the network as part of ospf instance id 3
`BR7(config-router)#exit`	Exit the router ospf mode
`BR7(config)#router bgp 100`	Configure router bgp
`BR7(config-router)#neighbor 3.3.3.3 remote-as 100`	Configure ibgp neighbor
`BR7(config-router)#neighbor 7.7.7.7 remote-as 100`	Configure ibgp neighbor
`BR7(config-router)#neighbor 8.8.8.8 remote-as 100`	Configure ibgp neighbor
`BR7(config-router)#neighbor 9.9.9.9 remote-as 100`	Configure ibgp neighbor
`BR7(config-router)#neighbor 13.13.13.13 remote-as 100`	Configure ibgp neighbor
`BR7(config-router)#neighbor 15.15.15.15 remote-as 100`	Configure ibgp neighbor
`BR7(config-router)#neighbor 3.3.3.3 update-source 12.12.12.12`	Configure update source for the neighbor with loopback address
`BR7(config-router)#neighbor 7.7.7.7 update-source 12.12.12.12`	Configure update source for the neighbor with loopback address
`BR7(config-router)#neighbor 8.8.8.8 update-source 12.12.12.12`	Configure update source for the neighbor with loopback address
`BR7(config-router)#neighbor 9.9.9.9 update-source 12.12.12.12`	Configure update source for the neighbor with loopback address
`BR7(config-router)#neighbor 13.13.13.13 update-source 12.12.12.12`	Configure update source for the neighbor with loopback address

© 2026 IP Infusion Inc. Proprietary

BR7(config-router)#neighbor 15.15.15.15 update-source 12.12.12.12	Configure update source for the neighbor with loopback address
BR7(config-router)#neighbor 3.3.3.3 fall-over bfd multihop	Configure bfd for the neighbor
BR7(config-router)#neighbor 7.7.7.7 fall-over bfd multihop	Configure bfd for the neighbor
BR7(config-router)#neighbor 8.8.8.8 fall-over bfd multihop	Configure bfd for the neighbor
BR7(config-router)#neighbor 9.9.9.9 fall-over bfd multihop	Configure bfd for the neighbor
BR7(config-router)#neighbor 13.13.13.13 fall-over bfd multihop	Configure bfd for the neighbor
BR7(config-router)#neighbor 15.15.15.15 fall-over bfd multihop	Configure bfd for the neighbor
BR7(config-router)#allocate-label all	Configure allocate label all command
BR7(config-router)#address-family ipv4 labeled-unicast	Configure address-family ipv4 labeled unicast
BR7(config-router-af)#neighbor 3.3.3.3 activate	Activate the neighbor under address family labeled unicast
BR7(config-router-af)#neighbor 7.7.7.7 activate	Activate the neighbor under address family labeled unicast
BR7(config-router-af)#neighbor 8.8.8.8 activate	Activate the neighbor under address family labeled unicast
BR7(config-router-af)#neighbor 9.9.9.9 activate	Activate the neighbor under address family labeled unicast
BR7(config-router-af)#neighbor 13.13.13.13 activate	Activate the neighbor under address family labeled unicast
BR7(config-router-af)#neighbor 15.15.15.15 activate	Activate the neighbor under address family labeled unicast
BR7(config-router-af)#neighbor 13.13.13.13 route-reflector-client	Configure neighbor as route reflector client
BR7(config-router-af)#neighbor 7.7.7.7 route-reflector-client	Configure neighbor as route reflector client
BR7(config-router-af)#neighbor 15.15.15.15 route-reflector-client	Configure neighbor as route reflector client
BR7(config-router-af)#neighbor 15.15.15.15 next-hop-self	Configure next-hop self for the neighbor
BR7(config-router-af)#neighbor 3.3.3.3 next-hop-self	Configure next-hop self for the neighbor
BR7(config-router-af)#neighbor 8.8.8.8 next-hop-self	Configure next-hop self for the neighbor
BR7(config-router-af)#neighbor 9.9.9.9 next-hop-self	Configure next-hop self for the neighbor
BR7(config-router-af)#bgp additional-paths send-receive	Enable additional paths send-receive capability on neighbor level under ipv4 labeled unicast address

© 2026 IP Infusion Inc. Proprietary

	family
BR7(config-router-af)#bgp additional-paths select all	Select all additional paths to advertise to the next router
BR7(config-router-af)#exit-address-family	Exit bgp address-family ipv6 labeled-unicast mode
BR7(config-router)#exit	Exit BGP router mode

P3 Configurations

P3(config)#interface lo	Enter interface lo mode
P3(config-if)#ip address 6.6.6.6/32	Configure ip address for loopback interface
P3(config-if)#exit	Exit interface loopback mode
P3(config)#router ldp	Configure ldp on the router
P3(config-router)#transport-address ipv4 6.6.6.60	Enable loopback ip as transport address
P3(config-router)#exit	Exit the router ldp mode
P3(config)#interface eth1	Enter interface mode
P3(config-if)#ip address 50.1.1.161/24	Configure ip address on the interface
P3(config-if)#enable-ldp ipv4	Enable ldp on the interface
P3(config-if)#label-switching	Enable label switching on the interface
P3(config-if)#exit	Exit the interface mode
P3(config)#interface eth2	Enter interface mode
P3(config-if)#ip address 60.1.1.161/24	Configure ip address on the interface
P3(config-if)#enable-ldp ipv4	Enable ldp on the interface
P3(config-if)#label-switching	Enable label switching on the interface
P3(config-if)#exit	Exit the interface mode
P3(config)#interface eth3	Enter interface mode
P3(config-if)#ip address 26.1.1.161/24	Configure ip address on the interface
P3(config-if)#enable-ldp ipv4	Enable rsvp on the interface
	Enable label switching on the interface
P3(config-if)#exit	Exit the interface mode
P3(config)#interface eth4	Enter interface mode
P3(config-if)#ip address 28.1.1.161/24	Configure ip address on the interface
P3(config-if)#enable-ldp ipv4	Enable ldp on the interface
P3(config-if)#label-switching	Enable label switching on the interface
P3(config-if)#exit	Exit the interface mode
P3(config)#interface eth5	Enter interface mode
P3(config-if)#ip address 33.1.1.161/24	Configure ip address on the interface

© 2026 IP Infusion Inc. Proprietary

`P3(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`P3(config-if)#label-switching`	Enable label switching on the interface
`P3(config-if)#exit`	Exit the interface mode
`P3(config)#interface eth6`	Enter interface mode
`P3(config-if)#ip address 32.1.1.161/24`	Configure ip address on the interface
`P3(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`P3(config-if)#label-switching`	Enable label switching on the interface
`P3(config-if)#exit`	Exit the interface mode
`P3(config)#interface eth7`	Enter interface mode
`P3(config-if)#ip address 42.1.1.161/24`	Configure ip address on the interface
`P3(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`P3(config-if)#label-switching`	Enable label switching on the interface
`P3(config-if)#exit`	Exit the interface mode
`P3(config)#router ospf 3`	Configure Router ospf with instance id 3
`P3(config-router)#network 26.1.1.0/24 area 0.0.0.0`	Add interface into ospf instance 3 with area 0 using network command
`P3(config-router)#network 28.1.1.0/24 area 0.0.0.0`	Add interface into ospf instance 3 with area 0 using network command
`P3(config-router)#network 32.1.1.0/24 area 0.0.0.0`	Add interface into ospf instance 3 with area 0 using network command
`P3(config-router)#network 33.1.1.0/24 area 0.0.0.0`	Add interface into ospf instance 3 with area 0 using network command
`P3(config-router)#network 42.1.1.0/24 area 0.0.0.0`	Add interface into ospf instance 3 with area 0 using network command
`P3(config-router)#network 50.1.1.0/24 area 0.0.0.0`	Add interface into ospf instance 3 with area 0 using network command
`P3(config-router)#network 60.1.1.0/24 area 0.0.0.0`	Add interface into ospf instance 3 with area 0 using network command
`P3(config-router)#network 6.6.6.6/32 area 0.0.0.0`	Add interface loopback into ospf instance 2 with area 0 using network command
`P3(config-router)#exit`	Exit router ospf mode

AGN4 Configurations

`AGN4(config)#interface lo`	Enter interface lo mode
`AGN4(config-if)#ip address 13.13.13.13/32`	Configure ip address for loopback interface
`AGN4(config-if)#exit`	Exit interface loopback mode
`AGN4(config)#router ldp`	Configure router ldp
`AGN4(config-router)#transport-address ipv4`	Configure transport address as loopback

© 2026 IP Infusion Inc. Proprietary

13.13.13.13	
AGN4(config)#exit	Exit the router ldp mode
AGN4(config)#interface eth6	Enter interface mode
AGN4(config-if)#ip address 32.1.1.170/24	Configure ip address on the interface
AGN4(config-if)#enable-ldp ipv4	Enable ldp on the interface
AGN4(config-if)#label-switching	Enable label-switching on the interface
AGN4(config-if)#exit	Exit the interface mode
AGN4(config)#router ospf 3	Configure router ospf instance 3
AGN4(config-router)#network 13.13.13.13/32 area 0	Add the loopback network into ospf instance 3
AGN4(config-router)#network 32.1.1.0/24 area 0	Add the ip address into ospf instance 3
AGN4(config-router)#exit	Exit router ospf mode
AGN4(config)#router bgp 100	Enter router bgp mode
AGN4(config-router)#network 13.13.13.13/32	Advertise loopback through network command inside bgp
AGN4(config-router)#neighbor 5.5.5.5 remote-as 100	Configure iBGP neighbor
AGN4(config-router)#neighbor 10.10.10.10 remote-as 100	Configure iBGP neighbor
AGN4(config-router)#neighbor 11.11.11.11 remote-as 100	Configure iBGP neighbor
AGN4(config-router)#neighbor 12.12.12.12 remote-as 100	Configure iBGP neighbor
AGN4(config-router)#neighbor 5.5.5.5 update-source 13.13.13.13	Update the source for the neighbor as loopback ip
AGN4(config-router)#neighbor 10.10.10.10 update-source 13.13.13.13	Update the source for the neighbor as loopback ip
AGN4(config-router)#neighbor 11.11.11.11 update-source 13.13.13.13	Update the source for the neighbor as loopback ip
AGN4(config-router)#neighbor 12.12.12.12 update-source 13.13.13.13	Update the source for the neighbor as loopback ip
AGN4(config-router)#neighbor 5.5.5.5 fall-over bfd multihop	Configure bfd for all the neighbors
AGN4(config-router)#neighbor 10.10.10.10 fall-over bfd multihop	Configure bfd for all the neighbors
AGN4(config-router)#neighbor 11.11.11.11 fall-over bfd multihop	Configure bfd for all the neighbors
AGN4(config-router)#neighbor 12.12.12.12 fall-over bfd multihop	Configure bfd for all the neighbors
AGN4(config-router)#allocate-label all	Enable allocate label all command
AGN4(config-router)#address-family ipv4 labeled-unicast	Enter ipv4 labeled-unicast address family
AGN4(config-router-af)#neighbor 5.5.5.5 activate	Activate the neighbor under ipv4 labeled-unicast address family

© 2026 IP Infusion Inc. Proprietary

`AGN4(config-router-af)#neighbor 10.10.10.10 activate`	Activate the neighbor under ipv4 labeled-unicast address family
`AGN4(config-router-af)#neighbor 11.11.11.11 activate`	Activate the neighbor under ipv4 labeled-unicast address family
`AGN4(config-router-af)#neighbor 12.12.12.12 activate`	Activate the neighbor under ipv4 labeled-unicast address family
`AGN4(config-router-af)#bgp additional-paths receive`	Enable bgp additional paths receive at global level under ipv4 labeled unicast address family
`AGN4(config-router-af)#exit`	End from router bgp mode

AGN2 Configurations

`AGN2(config)#interface lo`	Enter interface lo mode
`AGN2(config-if)#ip address 7.7.7.7/32`	Configure ip address for loopback interface
`AGN2(config-if)#exit`	Exit interface loopback mode
`AGN2(config)#router ldp`	Configure router ldp
`AGN2(config-router)#transport-address ipv4 7.7.7.7`	Configure transport address as loopback
`AGN2(config-router)#exit`	Exit router ldp mode
`AGN2(config)#interface eth2`	Enter interface mode
`AGN2(config-if)#ip address 60.1.1.162/24`	Configure ip address on the interface
`AGN2(config-if)#enable-ldp ipv4`	Enable ldp on the interface
`AGN2(config-if)#label-switching`	Enable label-switching on the interface
`AGN2(config-if)#exit`	Exit the interface mode
`AGN2(config)#router ospf 3`	Configure router ospf instance 3
`AGN2(config-router)#network 7.7.7.7/32 area 0`	Add the loopback network into ospf instance 3
`AGN2(config-router)#network 60.1.1.0/24 area 0`	Add the ip address into ospf instance 3
`AGN2(config-router)#exit`	Exit router ospf mode
`AGN2(config)#router bgp 100`	Enter router bgp mode
`AGN2(config-router)#network 7.7.7.7/32`	Advertise loopback through network command inside bgp
`AGN2(config-router)#neighbor 5.5.5.5 remote-as 100`	Configure iBGP neighbor
`AGN2(config-router)#neighbor 10.10.10.10 remote-as 100`	Configure iBGP neighbor
`AGN2(config-router)#neighbor 11.11.11.11 remote-as 100`	Configure iBGP neighbor
`AGN2(config-router)#neighbor 12.12.12.12 remote-as 100`	Configure iBGP neighbor
`AGN2(config-router)#neighbor 5.5.5.5 update-source 7.7.7.7`	Update the source for the neighbor as loopback ip

© 2026 IP Infusion Inc. Proprietary

`AGN2(config-router)#neighbor 10.10.10.10 update-source 7.7.7.7`	Update the source for the neighbor as loopback ip
`AGN2(config-router)#neighbor 11.11.11.11 update-source 7.7.7.7`	Update the source for the neighbor as loopback ip
`AGN2(config-router)#neighbor 12.12.12.12 update-source 7.7.7.7`	Update the source for the neighbor as loopback ip
`AGN2(config-router)#neighbor 5.5.5.5 fall-over bfd multihop`	Configure bfd for all the neighbors
`AGN2(config-router)#neighbor 10.10.10.10 fall-over bfd multihop`	Configure bfd for all the neighbors
`AGN2(config-router)#neighbor 11.11.11.11 fall-over bfd multihop`	Configure bfd for all the neighbors
`AGN2(config-router)#neighbor 12.12.12.12 fall-over bfd multihop`	Configure bfd for all the neighbors
`AGN2(config-router)#allocate-label all`	Enable allocate label all command
`AGN2(config-router)#address-family ipv4 labeled-unicast`	Enter ipv4 labeled-unicast address family
`AGN2(config-router-af)#neighbor 5.5.5.5 activate`	Activate the neighbor under ipv4 labeled-unicast address family
`AGN2(config-router-af)#neighbor 10.10.10.10 activate`	Activate the neighbor under ipv4 labeled-unicast address family
`AGN2(config-router-af)#neighbor 11.11.11.11 activate`	Activate the neighbor under ipv4 labeled-unicast address family
`AGN2(config-router-af)#neighbor 12.12.12.12 activate`	Activate the neighbor under ipv4 labeled-unicast address family
`AGN2(config-router-af)#neighbor 5.5.5.5 additional-paths receive`	Enable bgp additional paths receive at neighbor level under ipv4 labeled unicast address family
`AGN2(config-router-af)#neighbor 10.10.10.10 additional-paths receive`	Enable bgp additional paths receive at neighbor level under ipv4 labeled unicast address family
`AGN2(config-router-af)#neighbor 11.11.11.11 additional-paths receive`	Enable bgp additional paths receive at neighbor level under ipv4 labeled unicast address family
`AGN2(config-router-af)#exit-address-family`	Exit BGP address-family ipv4-labeled unicast mode
`AGN2(config-router)#end`	Return to configure terminal mode

AGN5 Configurations

`AGN5(config)#interface lo`	Enter interface lo mode
`AGN5(config-if)#ip address 15.15.15.15/32`	Configure ip address for loopback interface
`AGN5(config-if)#exit`	Exit interface loopback mode
`AGN5(config)#router ldp`	Configure router ldp
`AGN5(config-router)#transport-address ipv4 15.15.15.15`	Configure transport address as loopback

© 2026 IP Infusion Inc. Proprietary

AGN5(config-router)#exit	Exit router ldp mode
AGN5(config)#interface eth7	Enter interface mode
AGN5(config-if)#ip address 42.1.1.172/24	Configure ip address on the interface
AGN5(config-if)#enable-ldp ipv4	Enable ldp on the interface
AGN5(config-if)#label-switching	Enable label-switching on the interface
AGN5(config-if)#exit	Exit the interface mode
AGN5(config)#router ospf 3	Configure router ospf instance 3
AGN5(config-router)#network 15.15.15.15/32 area 0	Add the loopback network into ospf instance 3
AGN5(config-router)#network 42.1.1.0/24 area 0	Add the ip address into ospf instance 3
AGN5(config-router)#exit	Exit router ospf mode
AGN5(config)#router bgp 100	Enter router bgp mode
AGN5(config-router)#network 15.15.15.15/32	Advertise loopback through network command inside bgp
AGN5(config-router)#neighbor 5.5.5.5 remote-as 100	Configure iBGP neighbor
AGN5(config-router)#neighbor 10.10.10.10 remote-as 100	Configure iBGP neighbor
AGN5(config-router)#neighbor 11.11.11.11 remote-as 100	Configure iBGP neighbor
AGN5(config-router)#neighbor 12.12.12.12 remote-as 100	Configure iBGP neighbor
AGN5(config-router)#neighbor 5.5.5.5 update-source 15.15.15.15	Update the source for the neighbor as loopback ip
AGN5(config-router)#neighbor 10.10.10.10 update-source 15.15.15.15	Update the source for the neighbor as loopback ip
AGN5(config-router)#neighbor 11.11.11.11 update-source 15.15.15.15	Update the source for the neighbor as loopback ip
AGN5(config-router)#neighbor 12.12.12.12 update-source 15.15.15.15	Update the source for the neighbor as loopback ip
AGN5(config-router)#neighbor 5.5.5.5 fall-over bfd multihop	Configure bfd for all the neighbors
AGN5(config-router)#neighbor 10.10.10.10 fall-over bfd multihop	Configure bfd for all the neighbors
AGN5(config-router)#neighbor 11.11.11.11 fall-over bfd multihop	Configure bfd for all the neighbors
AGN5(config-router)#neighbor 12.12.12.12 fall-over bfd multihop	Configure bfd for all the neighbors
AGN5(config-router)#allocate-label all	Enable allocate label all command
AGN5(config-router)#address-family ipv4 labeled-unicast	Enter ipv4 labeled-unicast address family
AGN5(config-router-af)#neighbor 5.5.5.5 activate	Activate the neighbor under ipv4 labeled-unicast address family
AGN5(config-router-af)#neighbor 10.10.10.10 activate	Activate the neighbor under ipv4 labeled-unicast

© 2026 IP Infusion Inc. Proprietary

	address family
`AGN5(config-router-af)#neighbor 11.11.11.11 activate`	Activate the neighbor under ipv4 labeled-unicast address family
`AGN5(config-router-af)#neighbor 12.12.12.12 activate`	Activate the neighbor under ipv4 labeled-unicast address family
`AGN5(config-router-af)#exit-address-family`	Exit BGP address-family ipv4-labeled unicast mode
`AGN5(config-router)#end`	Return to configure terminal mode
`AGN5(config-router-af)#exit`	End from router bgp mode

Validation for BGP Labeled Unicast with Additional Paths

Verify OSPF neighborship is up with `show ip ospf neighbor`.

Verify LDP sessions are up with `show ldp session`.

Check BGPLU with Additional Paths

```
AGN1#

BR3#show ip bgp labeled-unicast all

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.1/32        1.1.1.1         24960(eth3)       24960
*>i 7.7.7.7/32        5.5.5.5            -           24960
*ai                   10.10.10.10        -           24962
* i                   11.11.11.11        -           24961
* i                   12.12.12.12        -           24962
*>i 13.13.13.13/32    10.10.10.10        -           24961
*ai                   5.5.5.5            -           24962
* i                   11.11.11.11        -           24962
* i                   12.12.12.12        -           24960
*>i 14.14.14.14/32    14.14.14.14     24961(eth3)       24960
*>i 15.15.15.15/32    5.5.5.5            -           24961
*ai                   10.10.10.10        -           24960
* i                   11.11.11.11        -           24960
* i                   12.12.12.12        -           24961
```

```
BR3#
BR3#show ip bgp labeled-unicast 13.13.13.13/32
BGP routing table entry for 13.13.13.13/32
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  14.14.14.14
  Local
    10.10.10.10 (metric 12) from 10.10.10.10 (10.12.28.167)
      Origin IGP, metric 0, localpref 100, label    24961, Out-label    24961, In-label  NA , valid,
internal, best
      Originator: 10.12.28.162, Cluster list: 10.12.28.167
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       1.1.1.1
      Last update: Mon Aug  6 00:10:36 2018

  Local
    5.5.5.5 (metric 12) from 5.5.5.5 (10.12.28.160)
      Origin IGP, metric 0, localpref 100, label    24962, Out-label    24962, In-label  NA , valid,
internal, backup, add-path
      Originator: 10.12.28.170, Cluster list: 10.12.28.160
      rx path_id: -1      tx path_id: 2
```

```
      Advertised to non peer-group peers:
       1.1.1.1
      Last update: Mon Aug  6 00:10:42 2018

  Local
    11.11.11.11 (metric 12) from 11.11.11.11 (10.12.28.168)
      Origin IGP, metric 0, localpref 100, label    24962, Out-label    24962, In-label  NA , valid,
internal
      Originator: 10.12.28.170, Cluster list: 10.12.28.168
      Not advertised to any peer
      Last update: Mon Aug  6 00:11:00 2018

  Local
    12.12.12.12 (metric 12) from 12.12.12.12 (10.12.28.169)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal
      Originator: 10.12.28.170, Cluster list: 10.12.28.169
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:29 2018
```

```
BGP PIC Enabled
BR3#show mpls ftn-table
 Primary FTN entry with FEC: 1.1.1.1/32, id: 2, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 14, in intf: - in label: 0 out-segment ix: 12
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 12, owner: LDP, Stale: NO, out intf: eth4, out label: 24342
    Nexthop addr: 21.1.1.157        cross connect ix: 14, op code: Push


 Primary FTN entry with FEC: 1.1.1.1/32, id: 10, row status: Active
  Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
      Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
       Out-segment with ix: 5, owner: BGP, Stale: NO, out intf: eth4, out label: 24960
    Nexthop addr: 1.1.1.1        cross connect ix: 5, op code: Push


 Primary FTN entry with FEC: 2.2.2.2/32, id: 3, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth4, out label: 3
    Nexthop addr: 21.1.1.157        cross connect ix: 2, op code: Push


 Primary FTN entry with FEC: 3.3.3.3/32, id: 4, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth4, out label: 24320
    Nexthop addr: 21.1.1.157        cross connect ix: 3, op code: Push


 Primary FTN entry with FEC: 4.4.4.4/32, id: 12, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 7
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth3, out label: 3
   Nexthop addr: 22.1.1.159        cross connect ix: 7, op code: Push


Primary FTN entry with FEC: 5.5.5.5/32, id: 13, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 8
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 8, owner: LDP, Stale: NO, out intf: eth3, out label: 24324
   Nexthop addr: 22.1.1.159        cross connect ix: 8, op code: Push


Primary FTN entry with FEC: 7.7.7.7/32, id: 24, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 16, in intf: - in label: 0 out-segment ix: 14
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 14, owner: BGP, Stale: NO, out intf: eth3, out label: 24960
   Nexthop addr: 5.5.5.5        cross connect ix: 16, op code: Push and Lookup


Non-primary FTN entry with FEC: 7.7.7.7/32, id: 26, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 18, in intf: - in label: 0 out-segment ix: 16
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 16, owner: BGP, Stale: NO, out intf: eth3, out label: 24962
   Nexthop addr: 10.10.10.10        cross connect ix: 18, op code: Push and Lookup


Primary FTN entry with FEC: 9.9.9.9/32, id: 5, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: eth4, out label: 24321
   Nexthop addr: 21.1.1.157        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 10.1.1.0/24, id: 6, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth4, out label: 3
   Nexthop addr: 21.1.1.157        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 10.10.10.10/32, id: 11, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth3, out label: 24329
   Nexthop addr: 22.1.1.159        cross connect ix: 6, op code: Push


Primary FTN entry with FEC: 11.11.11.11/32, id: 14, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
```

```
 Primary:  Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 10
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 10, owner: LDP, Stale: NO, out intf: eth3, out label: 24339
   Nexthop addr: 22.1.1.159        cross connect ix: 11, op code: Push


Primary FTN entry with FEC: 12.12.12.12/32, id: 15, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 9
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 9, owner: LDP, Stale: NO, out intf: eth3, out label: 24321
   Nexthop addr: 22.1.1.159        cross connect ix: 9, op code: Push


Primary FTN entry with FEC: 13.13.13.13/32, id: 22, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 12, in intf: - in label: 0 out-segment ix: 11
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 11, owner: BGP, Stale: NO, out intf: eth3, out label: 24961
   Nexthop addr: 10.10.10.10        cross connect ix: 12, op code: Push and Lookup


Non-primary FTN entry with FEC: 13.13.13.13/32, id: 28, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 20, in intf: - in label: 0 out-segment ix: 18
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 18, owner: BGP, Stale: NO, out intf: eth3, out label: 24962
   Nexthop addr: 5.5.5.5        cross connect ix: 20, op code: Push and Lookup


Primary FTN entry with FEC: 14.14.14.14/32, id: 1, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth4, out label: 24324
   Nexthop addr: 21.1.1.157        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 14.14.14.14/32, id: 27, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 19, in intf: - in label: 0 out-segment ix: 17
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 17, owner: BGP, Stale: NO, out intf: eth4, out label: 24960
   Nexthop addr: 14.14.14.14        cross connect ix: 19, op code: Push


Primary FTN entry with FEC: 15.15.15.15/32, id: 23, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 15, in intf: - in label: 0 out-segment ix: 13
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 13, owner: BGP, Stale: NO, out intf: eth3, out label: 24961
   Nexthop addr: 5.5.5.5        cross connect ix: 15, op code: Push and Lookup


Non-primary FTN entry with FEC: 15.15.15.15/32, id: 25, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
```

```
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 17, in intf: - in label: 0 out-segment ix: 15
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 15, owner: BGP, Stale: NO, out intf: eth3, out label: 24960
   Nexthop addr: 10.10.10.10        cross connect ix: 17, op code: Push and Lookup


Primary FTN entry with FEC: 20.1.1.0/24, id: 7, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth4, out label: 3
   Nexthop addr: 21.1.1.157        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 23.1.1.0/24, id: 8, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth4, out label: 3
   Nexthop addr: 21.1.1.157        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 24.1.1.0/24, id: 16, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 7
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth3, out label: 3
   Nexthop addr: 22.1.1.159        cross connect ix: 7, op code: Push


Primary FTN entry with FEC: 25.1.1.0/24, id: 17, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 7
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth3, out label: 3
   Nexthop addr: 22.1.1.159        cross connect ix: 7, op code: Push


Primary FTN entry with FEC: 27.1.1.0/24, id: 18, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 7
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth3, out label: 3
   Nexthop addr: 22.1.1.159        cross connect ix: 7, op code: Push


Primary FTN entry with FEC: 29.1.1.0/24, id: 19, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 7
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth3, out label: 3
   Nexthop addr: 22.1.1.159        cross connect ix: 7, op code: Push
```

```
 Primary FTN entry with FEC: 30.1.1.0/24, id: 20, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 7
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth3, out label: 3
    Nexthop addr: 22.1.1.159        cross connect ix: 7, op code: Push


 Primary FTN entry with FEC: 40.1.1.0/24, id: 21, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 7
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth3, out label: 3
    Nexthop addr: 22.1.1.159        cross connect ix: 7, op code: Push


 Primary FTN entry with FEC: 41.1.1.0/24, id: 9, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth4, out label: 3
    Nexthop addr: 21.1.1.157        cross connect ix: 2, op code: Push
```

```
BR3#

BR4#show ip bgp labeled-unicast all

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.1/32        1.1.1.1         24961(eth4)       24960
*>i 7.7.7.7/32        5.5.5.5            -           24960
*ai                   10.10.10.10        -           24962
*ai                   11.11.11.11        -           24961
*ai                   12.12.12.12        -           24962
*>i 13.13.13.13/32    10.10.10.10        -           24961
*ai                   5.5.5.5            -           24962
*ai                   11.11.11.11        -           24962
*ai                   12.12.12.12        -           24960
*>i 14.14.14.14/32    14.14.14.14     24960(eth4)       24960
*>i 15.15.15.15/32    5.5.5.5            -           24961
*ai                   10.10.10.10        -           24960
*ai                   11.11.11.11        -           24960
*ai                   12.12.12.12        -           24961
```

```
BR4#
BR4#show ip bgp labeled-unicast all summary
BGP router identifier 10.12.28.166, local AS number 100
BGP table version is 8
1 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  576        583       8      0      0  04:46:35              1
5.5.5.5                  4   100  577        576       8      0      0  04:46:34              3
10.10.10.10              4   100  575        575       8      0      0  04:46:21              3
11.11.11.11              4   100  577        576       8      0      0  04:46:31              3
12.12.12.12              4   100  578        577       8      0      0  04:46:50              3
14.14.14.14              4   100  576        581       8      0      0  04:46:52              1

Total number of neighbors 6
```

```
Total number of Established sessions 6
BR4#
BR4#show ip bgp labeled-unicast 13.13.13.13/32
BGP routing table entry for 13.13.13.13/32
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  14.14.14.14
  Local
    10.10.10.10 (metric 12) from 10.10.10.10 (10.12.28.167)
      Origin IGP, metric 0, localpref 100, label    24961, Out-label    24961, In-label  NA , valid,
internal, best
      Originator: 10.12.28.162, Cluster list: 10.12.28.167
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       1.1.1.1
      Last update: Mon Aug  6 00:10:50 2018

  Local
    5.5.5.5 (metric 12) from 5.5.5.5 (10.12.28.160)
      Origin IGP, metric 0, localpref 100, label    24962, Out-label    24962, In-label  NA , valid,
internal, backup, add-path
      Originator: 10.12.28.170, Cluster list: 10.12.28.160
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       1.1.1.1
      Last update: Mon Aug  6 00:10:40 2018

  Local
    11.11.11.11 (metric 12) from 11.11.11.11 (10.12.28.168)
      Origin IGP, metric 0, localpref 100, label    24962, Out-label    24962, In-label  NA , valid,
internal, add-path
      Originator: 10.12.28.170, Cluster list: 10.12.28.168
      rx path_id: -1      tx path_id: 3
      Advertised to non peer-group peers:
       1.1.1.1
      Last update: Mon Aug  6 00:11:02 2018

  Local
    12.12.12.12 (metric 12) from 12.12.12.12 (10.12.28.169)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal, add-path
      Originator: 10.12.28.170, Cluster list: 10.12.28.169
      rx path_id: -1      tx path_id: 0
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:28 2018
```

```
BGP PIC Enabled
BR4#

BR1#show ip bgp labeled-unicast all

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.1/32       1.1.1.1         24960(eth1)       24960
*>i 7.7.7.7/32       5.5.5.5            -           24960
*ai                  11.11.11.11        -           24961
*ai                  12.12.12.12        -           24962
*ai                  10.10.10.10        -           24962
*>i 13.13.13.13/32   5.5.5.5            -           24962
*ai                  10.10.10.10        -           24961
*ai                  11.11.11.11        -           24962
*ai                  12.12.12.12        -           24960
*>i 14.14.14.14/32   14.14.14.14     24961(eth1)       24960
*>i 15.15.15.15/32   5.5.5.5            -           24961
*ai                  12.12.12.12        -           24961
*ai                  10.10.10.10        -           24960
*ai                  11.11.11.11        -           24960
```

© 2026 IP Infusion Inc. Proprietary

```
BR1#
BR1#show ip bgp labeled-unicast all summary
BGP router identifier 10.12.28.158, local AS number 100
BGP table version is 7
1 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  578        587       7      0      0  04:47:48              1
5.5.5.5                  4   100  579        579       7      0      0  04:47:46              3
10.10.10.10              4   100  579        579       7      0      0  04:47:50              3
11.11.11.11              4   100  580        579       7      0      0  04:47:56              3
12.12.12.12              4   100  579        579       7      0      0  04:47:55              3
14.14.14.14              4   100  578        592       7      0      0  04:47:43              1

Total number of neighbors 6

Total number of Established sessions 6
BR1#
BR1#show ip bgp labeled-unicast 13.13.13.13/32
BGP routing table entry for 13.13.13.13/32
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  Local
    5.5.5.5 (metric 12) from 5.5.5.5 (10.12.28.160)
      Origin IGP, metric 0, localpref 100, label    24962, Out-label    24962, In-label  NA , valid,
internal, best
      Originator: 10.12.28.170, Cluster list: 10.12.28.160
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       1.1.1.1       14.14.14.14
      Last update: Mon Aug  6 00:10:43 2018

  Local
    10.10.10.10 (metric 12) from 10.10.10.10 (10.12.28.167)
      Origin IGP, metric 0, localpref 100, label    24961, Out-label    24961, In-label  NA , valid,
internal, backup, add-path
      Originator: 10.12.28.170, Cluster list: 10.12.28.167
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       1.1.1.1
      Last update: Mon Aug  6 00:10:34 2018

  Local
    11.11.11.11 (metric 12) from 11.11.11.11 (10.12.28.168)
      Origin IGP, metric 0, localpref 100, label    24962, Out-label    24962, In-label  NA , valid,
internal, add-path
      Originator: 10.12.28.170, Cluster list: 10.12.28.168
      rx path_id: -1      tx path_id: 3
      Advertised to non peer-group peers:
       1.1.1.1
      Last update: Mon Aug  6 00:11:00 2018

  Local
    12.12.12.12 (metric 12) from 12.12.12.12 (10.12.28.169)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal, add-path
      Originator: 10.12.28.170, Cluster list: 10.12.28.169
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       1.1.1.1
      Last update: Mon Aug  6 00:10:30 2018
```

```
BGP PIC Enabled
BR1#
BR1#show mpls ftn-table
 Primary FTN entry with FEC: 1.1.1.1/32, id: 1, row status: Active
```

```
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 18, in intf: - in label: 0 out-segment ix: 13
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 13, owner: LDP, Stale: NO, out intf: eth2, out label: 24340
   Nexthop addr: 20.1.1.157        cross connect ix: 18, op code: Push


Primary FTN entry with FEC: 1.1.1.1/32, id: 22, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 15, in intf: - in label: 0 out-segment ix: 11
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 11, owner: BGP, Stale: NO, out intf: eth2, out label: 24960
   Nexthop addr: 1.1.1.1        cross connect ix: 15, op code: Push


Primary FTN entry with FEC: 2.2.2.2/32, id: 2, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth2, out label: 3
   Nexthop addr: 20.1.1.157        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 4.4.4.4/32, id: 10, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth1, out label: 3
   Nexthop addr: 30.1.1.159        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 5.5.5.5/32, id: 11, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 7
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth1, out label: 24324
   Nexthop addr: 30.1.1.159        cross connect ix: 9, op code: Push


Primary FTN entry with FEC: 7.7.7.7/32, id: 25, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 20, in intf: - in label: 0 out-segment ix: 15
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 15, owner: BGP, Stale: NO, out intf: eth1, out label: 24960
   Nexthop addr: 5.5.5.5        cross connect ix: 20, op code: Push and Lookup


Non-primary FTN entry with FEC: 7.7.7.7/32, id: 26, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 21, in intf: - in label: 0 out-segment ix: 16
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 16, owner: BGP, Stale: NO, out intf: eth1, out label: 24961
   Nexthop addr: 11.11.11.11        cross connect ix: 21, op code: Push and Lookup
```

© 2026 IP Infusion Inc. Proprietary

```
Primary FTN entry with FEC: 8.8.8.8/32, id: 3, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: eth2, out label: 24325
   Nexthop addr: 20.1.1.157        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 9.9.9.9/32, id: 4, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth2, out label: 24321
   Nexthop addr: 20.1.1.157        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 10.1.1.0/24, id: 5, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth2, out label: 3
   Nexthop addr: 20.1.1.157        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 10.10.10.10/32, id: 12, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 8
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 8, owner: LDP, Stale: NO, out intf: eth1, out label: 24327
   Nexthop addr: 30.1.1.159        cross connect ix: 11, op code: Push


Primary FTN entry with FEC: 11.11.11.11/32, id: 13, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 9
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 9, owner: LDP, Stale: NO, out intf: eth1, out label: 24337
   Nexthop addr: 30.1.1.159        cross connect ix: 13, op code: Push


Primary FTN entry with FEC: 12.12.12.12/32, id: 14, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth1, out label: 24321
   Nexthop addr: 30.1.1.159        cross connect ix: 7, op code: Push


Primary FTN entry with FEC: 13.13.13.13/32, id: 24, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 19, in intf: - in label: 0 out-segment ix: 14
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 14, owner: BGP, Stale: NO, out intf: eth1, out label: 24962
```

© 2026 IP Infusion Inc. Proprietary

```
   Nexthop addr: 5.5.5.5        cross connect ix: 19, op code: Push and Lookup


Non-primary FTN entry with FEC: 13.13.13.13/32, id: 27, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 22, in intf: - in label: 0 out-segment ix: 17
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 17, owner: BGP, Stale: NO, out intf: eth1, out label: 24961
   Nexthop addr: 10.10.10.10        cross connect ix: 22, op code: Push and Lookup


Primary FTN entry with FEC: 14.14.14.14/32, id: 6, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth2, out label: 24322
   Nexthop addr: 20.1.1.157        cross connect ix: 3, op code: Push


Primary FTN entry with FEC: 14.14.14.14/32, id: 28, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 23, in intf: - in label: 0 out-segment ix: 18
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 18, owner: BGP, Stale: NO, out intf: eth2, out label: 24960
   Nexthop addr: 14.14.14.14        cross connect ix: 23, op code: Push


Primary FTN entry with FEC: 15.15.15.15/32, id: 21, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 14, in intf: - in label: 0 out-segment ix: 10
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 10, owner: BGP, Stale: NO, out intf: eth1, out label: 24961
   Nexthop addr: 5.5.5.5        cross connect ix: 14, op code: Push and Lookup


Non-primary FTN entry with FEC: 15.15.15.15/32, id: 23, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 16, in intf: - in label: 0 out-segment ix: 12
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 12, owner: BGP, Stale: NO, out intf: eth1, out label: 24961
   Nexthop addr: 12.12.12.12        cross connect ix: 16, op code: Push and Lookup


Primary FTN entry with FEC: 21.1.1.0/24, id: 7, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth2, out label: 3
   Nexthop addr: 20.1.1.157        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 22.1.1.0/24, id: 15, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
```

© 2026 IP Infusion Inc. Proprietary

```
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth1, out label: 3
   Nexthop addr: 30.1.1.159        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 23.1.1.0/24, id: 8, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth2, out label: 3
   Nexthop addr: 20.1.1.157        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 24.1.1.0/24, id: 16, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth1, out label: 3
   Nexthop addr: 30.1.1.159        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 25.1.1.0/24, id: 17, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth1, out label: 3
   Nexthop addr: 30.1.1.159        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 27.1.1.0/24, id: 18, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth1, out label: 3
   Nexthop addr: 30.1.1.159        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 29.1.1.0/24, id: 19, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth1, out label: 3
   Nexthop addr: 30.1.1.159        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 40.1.1.0/24, id: 20, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth1, out label: 3
   Nexthop addr: 30.1.1.159        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 41.1.1.0/24, id: 9, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
```

```
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth2, out label: 3
    Nexthop addr: 20.1.1.157        cross connect ix: 1, op code: Push

    BR1#
```

```
BR4#show ip bgp labeled-unicast all

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.1/32        1.1.1.1         24961(eth4)       24960
*>i 7.7.7.7/32        5.5.5.5            -           24960
*ai                   10.10.10.10        -           24962
*ai                   11.11.11.11        -           24961
*ai                   12.12.12.12        -           24962
*>i 13.13.13.13/32    10.10.10.10        -           24961
*ai                   5.5.5.5            -           24962
*ai                   11.11.11.11        -           24962
*ai                   12.12.12.12        -           24960
*>i 14.14.14.14/32    14.14.14.14     24960(eth4)       24960
*>i 15.15.15.15/32    5.5.5.5            -           24961
*ai                   10.10.10.10        -           24960
*ai                   11.11.11.11        -           24960
*ai                   12.12.12.12        -           24961
BR4#
BR4#show ip bgp labeled-unicast all summary
BGP router identifier 10.12.28.166, local AS number 100
BGP table version is 8
1 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  576        583       8      0      0  04:46:35               1
5.5.5.5                  4   100  577        576       8      0      0  04:46:34               3
10.10.10.10              4   100  575        575       8      0      0  04:46:21               3
11.11.11.11              4   100  577        576       8      0      0  04:46:31               3
12.12.12.12              4   100  578        577       8      0      0  04:46:50               3
14.14.14.14              4   100  576        581       8      0      0  04:46:52               1

Total number of neighbors 6

Total number of Established sessions 6
BR4#
BR4#show ip bgp labeled-unicast 13.13.13.13/32
BGP routing table entry for 13.13.13.13/32
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  14.14.14.14
  Local
    10.10.10.10 (metric 12) from 10.10.10.10 (10.12.28.167)
      Origin IGP, metric 0, localpref 100, label    24961, Out-label    24961, In-label  NA , valid,
internal, best
      Originator: 10.12.28.162, Cluster list: 10.12.28.167
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       1.1.1.1
      Last update: Mon Aug  6 00:10:50 2018

  Local
    5.5.5.5 (metric 12) from 5.5.5.5 (10.12.28.160)
      Origin IGP, metric 0, localpref 100, label    24962, Out-label    24962, In-label  NA , valid,
internal, backup, add-path
      Originator: 10.12.28.170, Cluster list: 10.12.28.160
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
```

```
       1.1.1.1
      Last update: Mon Aug  6 00:10:40 2018

  Local
    11.11.11.11 (metric 12) from 11.11.11.11 (10.12.28.168)
      Origin IGP, metric 0, localpref 100, label    24962, Out-label    24962, In-label  NA , valid,
internal, add-path
      Originator: 10.12.28.170, Cluster list: 10.12.28.168
      rx path_id: -1      tx path_id: 3
      Advertised to non peer-group peers:
       1.1.1.1
      Last update: Mon Aug  6 00:11:02 2018

  Local
    12.12.12.12 (metric 12) from 12.12.12.12 (10.12.28.169)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal, add-path
      Originator: 10.12.28.170, Cluster list: 10.12.28.169
      rx path_id: -1      tx path_id: 0
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:28 2018


BGP PIC Enabled
BR4#show  mpls ftn-table
 Primary FTN entry with FEC: 1.1.1.1/32, id: 6, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 16, in intf: - in label: 0 out-segment ix: 12
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 12, owner: LDP, Stale: NO, out intf: eth3, out label: 24341
    Nexthop addr: 23.1.1.157        cross connect ix: 16, op code: Push


 Primary FTN entry with FEC: 1.1.1.1/32, id: 24, row status: Active
  Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 18, in intf: - in label: 0 out-segment ix: 14
      Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
       Out-segment with ix: 14, owner: BGP, Stale: NO, out intf: eth3, out label: 24960
    Nexthop addr: 1.1.1.1        cross connect ix: 18, op code: Push


 Primary FTN entry with FEC: 2.2.2.2/32, id: 1, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth3, out label: 3
    Nexthop addr: 23.1.1.157        cross connect ix: 1, op code: Push


 Primary FTN entry with FEC: 3.3.3.3/32, id: 7, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth3, out label: 24320
    Nexthop addr: 23.1.1.157        cross connect ix: 2, op code: Push


 Primary FTN entry with FEC: 4.4.4.4/32, id: 11, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
```

```
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth4, out label: 3
   Nexthop addr: 24.1.1.159        cross connect ix: 6, op code: Push


Primary FTN entry with FEC: 5.5.5.5/32, id: 12, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 8
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 8, owner: LDP, Stale: NO, out intf: eth4, out label: 24325
   Nexthop addr: 24.1.1.159        cross connect ix: 9, op code: Push


Primary FTN entry with FEC: 7.7.7.7/32, id: 25, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 19, in intf: - in label: 0 out-segment ix: 15
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 15, owner: BGP, Stale: NO, out intf: eth4, out label: 24960
   Nexthop addr: 5.5.5.5        cross connect ix: 19, op code: Push and Lookup


Non-primary FTN entry with FEC: 7.7.7.7/32, id: 27, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 21, in intf: - in label: 0 out-segment ix: 17
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 17, owner: BGP, Stale: NO, out intf: eth4, out label: 24962
   Nexthop addr: 10.10.10.10        cross connect ix: 21, op code: Push and Lookup


Primary FTN entry with FEC: 8.8.8.8/32, id: 8, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: eth3, out label: 24326
   Nexthop addr: 23.1.1.157        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 10.1.1.0/24, id: 2, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth3, out label: 3
   Nexthop addr: 23.1.1.157        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 10.10.10.10/32, id: 13, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 9
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 9, owner: LDP, Stale: NO, out intf: eth4, out label: 24330
   Nexthop addr: 24.1.1.159        cross connect ix: 11, op code: Push


Primary FTN entry with FEC: 11.11.11.11/32, id: 14, row status: Active
```

```
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 10
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 10, owner: LDP, Stale: NO, out intf: eth4, out label: 24340
   Nexthop addr: 24.1.1.159        cross connect ix: 13, op code: Push


Primary FTN entry with FEC: 12.12.12.12/32, id: 15, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 7
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth4, out label: 24321
   Nexthop addr: 24.1.1.159        cross connect ix: 7, op code: Push


Primary FTN entry with FEC: 13.13.13.13/32, id: 22, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 14, in intf: - in label: 0 out-segment ix: 11
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 11, owner: BGP, Stale: NO, out intf: eth4, out label: 24961
   Nexthop addr: 10.10.10.10        cross connect ix: 14, op code: Push and Lookup


Non-primary FTN entry with FEC: 13.13.13.13/32, id: 28, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 22, in intf: - in label: 0 out-segment ix: 18
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 18, owner: BGP, Stale: NO, out intf: eth4, out label: 24962
   Nexthop addr: 5.5.5.5        cross connect ix: 22, op code: Push and Lookup


Primary FTN entry with FEC: 14.14.14.14/32, id: 9, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth3, out label: 24323
   Nexthop addr: 23.1.1.157        cross connect ix: 3, op code: Push


Primary FTN entry with FEC: 14.14.14.14/32, id: 10, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 5, owner: BGP, Stale: NO, out intf: eth3, out label: 24960
   Nexthop addr: 14.14.14.14        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 15.15.15.15/32, id: 23, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 17, in intf: - in label: 0 out-segment ix: 13
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 13, owner: BGP, Stale: NO, out intf: eth4, out label: 24961
   Nexthop addr: 5.5.5.5        cross connect ix: 17, op code: Push and Lookup
```

```
Non-primary FTN entry with FEC: 15.15.15.15/32, id: 26, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 20, in intf: - in label: 0 out-segment ix: 16
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 16, owner: BGP, Stale: NO, out intf: eth4, out label: 24960
   Nexthop addr: 10.10.10.10        cross connect ix: 20, op code: Push and Lookup


Primary FTN entry with FEC: 20.1.1.0/24, id: 3, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth3, out label: 3
   Nexthop addr: 23.1.1.157        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 21.1.1.0/24, id: 4, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth3, out label: 3
   Nexthop addr: 23.1.1.157        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 22.1.1.0/24, id: 16, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth4, out label: 3
   Nexthop addr: 24.1.1.159        cross connect ix: 6, op code: Push


Primary FTN entry with FEC: 25.1.1.0/24, id: 17, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth4, out label: 3
   Nexthop addr: 24.1.1.159        cross connect ix: 6, op code: Push


Primary FTN entry with FEC: 27.1.1.0/24, id: 18, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth4, out label: 3
   Nexthop addr: 24.1.1.159        cross connect ix: 6, op code: Push


Primary FTN entry with FEC: 29.1.1.0/24, id: 19, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth4, out label: 3
```

© 2026 IP Infusion Inc. Proprietary

```
    Nexthop addr: 24.1.1.159        cross connect ix: 6, op code: Push


 Primary FTN entry with FEC: 30.1.1.0/24, id: 20, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth4, out label: 3
    Nexthop addr: 24.1.1.159        cross connect ix: 6, op code: Push


 Primary FTN entry with FEC: 40.1.1.0/24, id: 21, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth4, out label: 3
    Nexthop addr: 24.1.1.159        cross connect ix: 6, op code: Push


 Primary FTN entry with FEC: 41.1.1.0/24, id: 5, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth3, out label: 3
    Nexthop addr: 23.1.1.157        cross connect ix: 1, op code: Push


BR4#
```

```
BR5#show ip bgp labeled-unicast all

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.1/32       3.3.3.3            -           24960
*ai                  8.8.8.8            -           24960
*ai                  9.9.9.9            -           24961
*>i 7.7.7.7/32       7.7.7.7         24962(eth5)       24960
*>i 13.13.13.13/32   13.13.13.13     24961(eth5)       24960
*>i 14.14.14.14/32   8.8.8.8            -           24961
*ai                  9.9.9.9            -           24960
*ai                  3.3.3.3            -           24961
*>i 15.15.15.15/32   15.15.15.15     24960(eth5)       24960
BR5#
BR5#show ip bgp labeled-unicast all summary
BGP router identifier 10.12.28.167, local AS number 100
BGP table version is 5
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
3.3.3.3                  4   100  587        587       5      0      0  04:51:40               2
7.7.7.7                  4   100  586        591       5      0      0  04:51:41               1
8.8.8.8                  4   100  585        585       5      0      0  04:51:30               2
9.9.9.9                  4   100  585        585       5      0      0  04:51:16               2
13.13.13.13              4   100  586        591       5      0      0  04:51:34               1
15.15.15.15              4   100  586        589       5      0      0  04:51:53               1

Total number of neighbors 6

Total number of Established sessions 6
BR5#
```

```
BR5#show ip bgp labeled-unicast 1.1.1.1/32
BGP routing table entry for 1.1.1.1/32
Paths: (3 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  15.15.15.15
  Local
    3.3.3.3 (metric 12) from 3.3.3.3 (10.12.28.158)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal, best
      Originator: 10.12.28.156, Cluster list: 10.12.28.158
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       7.7.7.7       13.13.13.13
      Last update: Mon Aug  6 00:10:30 2018

  Local
    8.8.8.8 (metric 12) from 8.8.8.8 (10.12.28.165)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal, add-path
      Originator: 10.12.28.156, Cluster list: 10.12.28.165
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       7.7.7.7       13.13.13.13
      Last update: Mon Aug  6 00:10:37 2018

  Local
    9.9.9.9 (metric 12) from 9.9.9.9 (10.12.28.166)
      Origin IGP, metric 0, localpref 100, label    24961, Out-label    24961, In-label  NA , valid,
internal, add-path
      Originator: 10.12.28.156, Cluster list: 10.12.28.166
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       7.7.7.7       13.13.13.13
      Last update: Mon Aug  6 00:10:51 2018

BR5#sh mpls ftn-table
 Primary FTN entry with FEC: 1.1.1.1/32, id: 25, row status: Active
  Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 18, in intf: - in label: 0 out-segment ix: 13
      Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 13, owner: BGP, Stale: NO, out intf: eth5, out label: 24960
    Nexthop addr: 3.3.3.3        cross connect ix: 18, op code: Push and Lookup


 Primary FTN entry with FEC: 3.3.3.3/32, id: 13, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: eth5, out label: 24320
    Nexthop addr: 25.1.1.159        cross connect ix: 4, op code: Push


 Primary FTN entry with FEC: 4.4.4.4/32, id: 15, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth5, out label: 3
    Nexthop addr: 25.1.1.159        cross connect ix: 5, op code: Push


 Primary FTN entry with FEC: 5.5.5.5/32, id: 10, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
```

```
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 10
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 10, owner: LDP, Stale: NO, out intf: eth3, out label: 24331
   Nexthop addr: 26.1.1.161        cross connect ix: 13, op code: Push


Primary FTN entry with FEC: 6.6.6.6/32, id: 1, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth3, out label: 3
   Nexthop addr: 26.1.1.161        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 7.7.7.7/32, id: 11, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 15, in intf: - in label: 0 out-segment ix: 11
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 11, owner: LDP, Stale: NO, out intf: eth3, out label: 24335
   Nexthop addr: 26.1.1.161        cross connect ix: 15, op code: Push


Primary FTN entry with FEC: 7.7.7.7/32, id: 27, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 20, in intf: - in label: 0 out-segment ix: 15
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 15, owner: BGP, Stale: NO, out intf: eth3, out label: 24960
   Nexthop addr: 7.7.7.7        cross connect ix: 20, op code: Push


Primary FTN entry with FEC: 8.8.8.8/32, id: 16, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 8
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 8, owner: LDP, Stale: NO, out intf: eth5, out label: 24335
   Nexthop addr: 25.1.1.159        cross connect ix: 9, op code: Push


Primary FTN entry with FEC: 9.9.9.9/32, id: 17, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth5, out label: 24322
   Nexthop addr: 25.1.1.159        cross connect ix: 6, op code: Push


Primary FTN entry with FEC: 11.11.11.11/32, id: 18, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 9
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 9, owner: LDP, Stale: NO, out intf: eth5, out label: 24341
   Nexthop addr: 25.1.1.159        cross connect ix: 11, op code: Push
```

© 2026 IP Infusion Inc. Proprietary

```
Primary FTN entry with FEC: 12.12.12.12/32, id: 2, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth3, out label: 24320
   Nexthop addr: 26.1.1.161        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 13.13.13.13/32, id: 12, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 17, in intf: - in label: 0 out-segment ix: 12
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 12, owner: LDP, Stale: NO, out intf: eth3, out label: 24340
   Nexthop addr: 26.1.1.161        cross connect ix: 17, op code: Push


Primary FTN entry with FEC: 13.13.13.13/32, id: 26, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 19, in intf: - in label: 0 out-segment ix: 14
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 14, owner: BGP, Stale: NO, out intf: eth3, out label: 24960
   Nexthop addr: 13.13.13.13        cross connect ix: 19, op code: Push


Primary FTN entry with FEC: 14.14.14.14/32, id: 28, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 21, in intf: - in label: 0 out-segment ix: 16
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 16, owner: BGP, Stale: NO, out intf: eth5, out label: 24961
   Nexthop addr: 8.8.8.8        cross connect ix: 21, op code: Push and Lookup


Primary FTN entry with FEC: 15.15.15.15/32, id: 3, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth3, out label: 24321
   Nexthop addr: 26.1.1.161        cross connect ix: 3, op code: Push


Primary FTN entry with FEC: 15.15.15.15/32, id: 14, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 7
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 7, owner: BGP, Stale: NO, out intf: eth3, out label: 24960
   Nexthop addr: 15.15.15.15        cross connect ix: 7, op code: Push


Primary FTN entry with FEC: 22.1.1.0/24, id: 19, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth5, out label: 3
   Nexthop addr: 25.1.1.159        cross connect ix: 5, op code: Push
```

© 2026 IP Infusion Inc. Proprietary

```
Primary FTN entry with FEC: 24.1.1.0/24, id: 20, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth5, out label: 3
   Nexthop addr: 25.1.1.159        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 27.1.1.0/24, id: 21, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth5, out label: 3
   Nexthop addr: 25.1.1.159        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 28.1.1.0/24, id: 4, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth3, out label: 3
   Nexthop addr: 26.1.1.161        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 29.1.1.0/24, id: 22, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth5, out label: 3
   Nexthop addr: 25.1.1.159        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 30.1.1.0/24, id: 23, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth5, out label: 3
   Nexthop addr: 25.1.1.159        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 32.1.1.0/24, id: 5, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth3, out label: 3
   Nexthop addr: 26.1.1.161        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 33.1.1.0/24, id: 6, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth3, out label: 3
    Nexthop addr: 26.1.1.161        cross connect ix: 1, op code: Push


 Primary FTN entry with FEC: 40.1.1.0/24, id: 24, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth5, out label: 3
    Nexthop addr: 25.1.1.159        cross connect ix: 5, op code: Push


 Primary FTN entry with FEC: 42.1.1.0/24, id: 7, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth3, out label: 3
    Nexthop addr: 26.1.1.161        cross connect ix: 1, op code: Push


 Primary FTN entry with FEC: 50.1.1.0/24, id: 8, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth3, out label: 3
    Nexthop addr: 26.1.1.161        cross connect ix: 1, op code: Push


 Primary FTN entry with FEC: 60.1.1.0/24, id: 9, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth3, out label: 3
    Nexthop addr: 26.1.1.161        cross connect ix: 1, op code: Push


BR5#
```

```
BR6#show ip bgp labeled-unicast all

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.1/32        3.3.3.3            -          24960
*ai                   8.8.8.8            -          24960
* i                   9.9.9.9            -          24961
*>i 7.7.7.7/32        7.7.7.7         24961(eth6)       24960
*>i 13.13.13.13/32    13.13.13.13     24962(eth6)       24960
*>i 14.14.14.14/32    9.9.9.9            -          24960
*ai                   3.3.3.3            -          24961
* i                   8.8.8.8            -          24961
*>i 15.15.15.15/32    15.15.15.15     24960(eth6)       24960
BR6#
BR6#
BR6#show ip bgp labeled-unicast all summary
BGP router identifier 10.12.28.168, local AS number 100
BGP table version is 6
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
```

```
3.3.3.3                 4   100  600        601       6      0      0  04:58:04              2
7.7.7.7                 4   100  598        603       6      0      0  04:57:57              1
8.8.8.8                 4   100  599        600       6      0      0  04:57:59              2
9.9.9.9                 4   100  598        599       6      0      0  04:57:45              2
13.13.13.13             4   100  597        602       6      0      0  04:57:29              1
15.15.15.15             4   100  598        601       6      0      0  04:57:57              1

Total number of neighbors 6

Total number of Established sessions 6
BR6#show ip bgp labeled-unicast 1.1.1.1/32
BGP routing table entry for 1.1.1.1/32
Paths: (3 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  15.15.15.15
  Local
    3.3.3.3 (metric 12) from 3.3.3.3 (10.12.28.158)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal, best
      Originator: 10.12.28.156, Cluster list: 10.12.28.158
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       7.7.7.7       13.13.13.13
      Last update: Mon Aug  6 00:10:29 2018

  Local
    8.8.8.8 (metric 12) from 8.8.8.8 (10.12.28.165)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal, add-path
      Originator: 10.12.28.156, Cluster list: 10.12.28.165
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       7.7.7.7       13.13.13.13
      Last update: Mon Aug  6 00:10:26 2018

  Local
    9.9.9.9 (metric 12) from 9.9.9.9 (10.12.28.166)
      Origin IGP, metric 0, localpref 100, label    24961, Out-label    24961, In-label  NA , valid,
internal
      Originator: 10.12.28.156, Cluster list: 10.12.28.166
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:40 2018

BR6#sh mpls ftn-table
 Primary FTN entry with FEC: 1.1.1.1/32, id: 25, row status: Active
  Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 14, in intf: - in label: 0 out-segment ix: 13
      Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 13, owner: BGP, Stale: NO, out intf: eth6, out label: 24960
    Nexthop addr: 3.3.3.3        cross connect ix: 14, op code: Push and Lookup


 Primary FTN entry with FEC: 3.3.3.3/32, id: 1, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth6, out label: 24320
    Nexthop addr: 27.1.1.159        cross connect ix: 1, op code: Push


 Primary FTN entry with FEC: 4.4.4.4/32, id: 2, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
```

```
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth6, out label: 3
   Nexthop addr: 27.1.1.159        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 5.5.5.5/32, id: 12, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth4, out label: 24331
   Nexthop addr: 28.1.1.161        cross connect ix: 7, op code: Push


Primary FTN entry with FEC: 6.6.6.6/32, id: 13, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 7
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth4, out label: 3
   Nexthop addr: 28.1.1.161        cross connect ix: 8, op code: Push


Primary FTN entry with FEC: 7.7.7.7/32, id: 14, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 8
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 8, owner: LDP, Stale: NO, out intf: eth4, out label: 24334
   Nexthop addr: 28.1.1.161        cross connect ix: 9, op code: Push


Primary FTN entry with FEC: 7.7.7.7/32, id: 26, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 15, in intf: - in label: 0 out-segment ix: 14
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 14, owner: BGP, Stale: NO, out intf: eth4, out label: 24960
   Nexthop addr: 7.7.7.7        cross connect ix: 15, op code: Push


Primary FTN entry with FEC: 8.8.8.8/32, id: 3, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth6, out label: 24332
   Nexthop addr: 27.1.1.159        cross connect ix: 3, op code: Push


Primary FTN entry with FEC: 9.9.9.9/32, id: 4, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: eth6, out label: 24322
   Nexthop addr: 27.1.1.159        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 10.10.10.10/32, id: 15, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
```

© 2026 IP Infusion Inc. Proprietary

```
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 9
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 9, owner: LDP, Stale: NO, out intf: eth4, out label: 24323
   Nexthop addr: 28.1.1.161        cross connect ix: 10, op code: Push


Primary FTN entry with FEC: 12.12.12.12/32, id: 16, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 10
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 10, owner: LDP, Stale: NO, out intf: eth4, out label: 24320
   Nexthop addr: 28.1.1.161        cross connect ix: 11, op code: Push


Primary FTN entry with FEC: 13.13.13.13/32, id: 11, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth4, out label: 24341
   Nexthop addr: 28.1.1.161        cross connect ix: 6, op code: Push


Primary FTN entry with FEC: 13.13.13.13/32, id: 28, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 17, in intf: - in label: 0 out-segment ix: 16
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 16, owner: BGP, Stale: NO, out intf: eth4, out label: 24960
   Nexthop addr: 13.13.13.13        cross connect ix: 17, op code: Push


Primary FTN entry with FEC: 14.14.14.14/32, id: 27, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 16, in intf: - in label: 0 out-segment ix: 15
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 15, owner: BGP, Stale: NO, out intf: eth6, out label: 24960
   Nexthop addr: 9.9.9.9        cross connect ix: 16, op code: Push and Lookup


Primary FTN entry with FEC: 15.15.15.15/32, id: 17, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 12, in intf: - in label: 0 out-segment ix: 11
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 11, owner: LDP, Stale: NO, out intf: eth4, out label: 24321
   Nexthop addr: 28.1.1.161        cross connect ix: 12, op code: Push


Primary FTN entry with FEC: 15.15.15.15/32, id: 24, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 12
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 12, owner: BGP, Stale: NO, out intf: eth4, out label: 24960
   Nexthop addr: 15.15.15.15        cross connect ix: 13, op code: Push
```

© 2026 IP Infusion Inc. Proprietary

```
Primary FTN entry with FEC: 22.1.1.0/24, id: 5, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth6, out label: 3
   Nexthop addr: 27.1.1.159        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 24.1.1.0/24, id: 6, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth6, out label: 3
   Nexthop addr: 27.1.1.159        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 25.1.1.0/24, id: 7, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth6, out label: 3
   Nexthop addr: 27.1.1.159        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 26.1.1.0/24, id: 18, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 7
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth4, out label: 3
   Nexthop addr: 28.1.1.161        cross connect ix: 8, op code: Push


Primary FTN entry with FEC: 29.1.1.0/24, id: 8, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth6, out label: 3
   Nexthop addr: 27.1.1.159        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 30.1.1.0/24, id: 9, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth6, out label: 3
   Nexthop addr: 27.1.1.159        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 32.1.1.0/24, id: 19, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 7
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth4, out label: 3
   Nexthop addr: 28.1.1.161        cross connect ix: 8, op code: Push
```

© 2026 IP Infusion Inc. Proprietary

```
 Primary FTN entry with FEC: 33.1.1.0/24, id: 20, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 7
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth4, out label: 3
    Nexthop addr: 28.1.1.161        cross connect ix: 8, op code: Push


 Primary FTN entry with FEC: 40.1.1.0/24, id: 10, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth6, out label: 3
    Nexthop addr: 27.1.1.159        cross connect ix: 2, op code: Push


 Primary FTN entry with FEC: 42.1.1.0/24, id: 21, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 7
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth4, out label: 3
    Nexthop addr: 28.1.1.161        cross connect ix: 8, op code: Push


 Primary FTN entry with FEC: 50.1.1.0/24, id: 22, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 7
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth4, out label: 3
    Nexthop addr: 28.1.1.161        cross connect ix: 8, op code: Push


 Primary FTN entry with FEC: 60.1.1.0/24, id: 23, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 7
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth4, out label: 3
    Nexthop addr: 28.1.1.161        cross connect ix: 8, op code: Push


BR6#
```

```
BR2#show ip bgp labeled-unicast all

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.1/32        3.3.3.3            -           24960
*ai                   8.8.8.8            -           24960
*ai                   9.9.9.9            -           24961
*>i 7.7.7.7/32        7.7.7.7         24960(eth2)       24960
*>i 13.13.13.13/32    13.13.13.13     24962(eth2)       24960
*>i 14.14.14.14/32    9.9.9.9            -           24960
*ai                   3.3.3.3            -           24961
*ai                   8.8.8.8            -           24961
*>i 15.15.15.15/32    15.15.15.15     24961(eth2)       24960
BR2#
```

```
BR2#show ip bgp labeled-unicast all summary
BGP router identifier 10.12.28.160, local AS number 100
BGP table version is 6
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
3.3.3.3                  4   100  601        601       6      0      0  04:58:39              2
7.7.7.7                  4   100  600        606       6      0      0  04:58:44              1
8.8.8.8                  4   100  601        601       6      0      0  04:58:43              2
9.9.9.9                  4   100  600        601       6      0      0  04:58:33              2
13.13.13.13              4   100  600        606       6      0      0  04:58:34              1
15.15.15.15              4   100  600        604       6      0      0  04:58:43              1

Total number of neighbors 6

Total number of Established sessions 6
BR2#
BR2#show ip bgp labeled-unicast 1.1.1.1/32
BGP routing table entry for 1.1.1.1/32
Paths: (3 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  15.15.15.15
  Local
    3.3.3.3 (metric 12) from 3.3.3.3 (10.12.28.158)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal, best
      Originator: 10.12.28.156, Cluster list: 10.12.28.158
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       7.7.7.7       13.13.13.13
      Last update: Mon Aug  6 00:10:31 2018

  Local
    8.8.8.8 (metric 12) from 8.8.8.8 (10.12.28.165)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal, add-path
      Originator: 10.12.28.156, Cluster list: 10.12.28.165
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       7.7.7.7       13.13.13.13
      Last update: Mon Aug  6 00:10:27 2018

  Local
    9.9.9.9 (metric 12) from 9.9.9.9 (10.12.28.166)
      Origin IGP, metric 0, localpref 100, label    24961, Out-label    24961, In-label  NA , valid,
internal, add-path
      Originator: 10.12.28.156, Cluster list: 10.12.28.166
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       7.7.7.7       13.13.13.13
      Last update: Mon Aug  6 00:10:37 2018

BR2#show mpls ftn-table
 Primary FTN entry with FEC: 1.1.1.1/32, id: 24, row status: Active
  Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 14, in intf: - in label: 0 out-segment ix: 12
      Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 12, owner: BGP, Stale: NO, out intf: eth2, out label: 24960
    Nexthop addr: 3.3.3.3        cross connect ix: 14, op code: Push and Lookup


 Primary FTN entry with FEC: 3.3.3.3/32, id: 1, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
```

```
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth2, out label: 24320
   Nexthop addr: 40.1.1.159        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 4.4.4.4/32, id: 2, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth2, out label: 3
   Nexthop addr: 40.1.1.159        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 6.6.6.6/32, id: 12, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth1, out label: 3
   Nexthop addr: 50.1.1.161        cross connect ix: 6, op code: Push


Primary FTN entry with FEC: 7.7.7.7/32, id: 13, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 10
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 10, owner: LDP, Stale: NO, out intf: eth1, out label: 24334
   Nexthop addr: 50.1.1.161        cross connect ix: 11, op code: Push


Primary FTN entry with FEC: 7.7.7.7/32, id: 25, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 15, in intf: - in label: 0 out-segment ix: 13
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 13, owner: BGP, Stale: NO, out intf: eth1, out label: 24960
   Nexthop addr: 7.7.7.7        cross connect ix: 15, op code: Push


Primary FTN entry with FEC: 8.8.8.8/32, id: 3, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: eth2, out label: 24333
   Nexthop addr: 40.1.1.159        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 9.9.9.9/32, id: 4, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth2, out label: 24322
   Nexthop addr: 40.1.1.159        cross connect ix: 3, op code: Push


Primary FTN entry with FEC: 10.10.10.10/32, id: 14, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
```

```
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 7
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth1, out label: 24323
   Nexthop addr: 50.1.1.161        cross connect ix: 7, op code: Push


Primary FTN entry with FEC: 11.11.11.11/32, id: 5, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth2, out label: 24338
   Nexthop addr: 40.1.1.159        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 12.12.12.12/32, id: 15, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 8
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 8, owner: LDP, Stale: NO, out intf: eth1, out label: 24320
   Nexthop addr: 50.1.1.161        cross connect ix: 8, op code: Push


Primary FTN entry with FEC: 13.13.13.13/32, id: 16, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 11
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 11, owner: LDP, Stale: NO, out intf: eth1, out label: 24338
   Nexthop addr: 50.1.1.161        cross connect ix: 13, op code: Push


Primary FTN entry with FEC: 13.13.13.13/32, id: 28, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 18, in intf: - in label: 0 out-segment ix: 16
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 16, owner: BGP, Stale: NO, out intf: eth1, out label: 24960
   Nexthop addr: 13.13.13.13        cross connect ix: 18, op code: Push


Primary FTN entry with FEC: 14.14.14.14/32, id: 27, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 17, in intf: - in label: 0 out-segment ix: 15
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 15, owner: BGP, Stale: NO, out intf: eth2, out label: 24960
   Nexthop addr: 9.9.9.9        cross connect ix: 17, op code: Push and Lookup


Primary FTN entry with FEC: 15.15.15.15/32, id: 17, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 9
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 9, owner: LDP, Stale: NO, out intf: eth1, out label: 24321
   Nexthop addr: 50.1.1.161        cross connect ix: 9, op code: Push
```

```
Primary FTN entry with FEC: 15.15.15.15/32, id: 26, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 16, in intf: - in label: 0 out-segment ix: 14
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 14, owner: BGP, Stale: NO, out intf: eth1, out label: 24960
   Nexthop addr: 15.15.15.15        cross connect ix: 16, op code: Push


Primary FTN entry with FEC: 22.1.1.0/24, id: 6, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth2, out label: 3
   Nexthop addr: 40.1.1.159        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 24.1.1.0/24, id: 7, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth2, out label: 3
   Nexthop addr: 40.1.1.159        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 25.1.1.0/24, id: 8, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth2, out label: 3
   Nexthop addr: 40.1.1.159        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 26.1.1.0/24, id: 18, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth1, out label: 3
   Nexthop addr: 50.1.1.161        cross connect ix: 6, op code: Push


Primary FTN entry with FEC: 27.1.1.0/24, id: 9, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth2, out label: 3
   Nexthop addr: 40.1.1.159        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 28.1.1.0/24, id: 19, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth1, out label: 3
   Nexthop addr: 50.1.1.161        cross connect ix: 6, op code: Push
```

© 2026 IP Infusion Inc. Proprietary

```
Primary FTN entry with FEC: 29.1.1.0/24, id: 10, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth2, out label: 3
   Nexthop addr: 40.1.1.159        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 30.1.1.0/24, id: 11, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth2, out label: 3
   Nexthop addr: 40.1.1.159        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 32.1.1.0/24, id: 20, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth1, out label: 3
   Nexthop addr: 50.1.1.161        cross connect ix: 6, op code: Push


Primary FTN entry with FEC: 33.1.1.0/24, id: 21, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth1, out label: 3
   Nexthop addr: 50.1.1.161        cross connect ix: 6, op code: Push


Primary FTN entry with FEC: 42.1.1.0/24, id: 22, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth1, out label: 3
   Nexthop addr: 50.1.1.161        cross connect ix: 6, op code: Push


Primary FTN entry with FEC: 60.1.1.0/24, id: 23, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth1, out label: 3
   Nexthop addr: 50.1.1.161        cross connect ix: 6, op code: Push


BR2#
```

```
BR7#show ip bgp labeled-unicast all

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
```

```
*>i 1.1.1.1/32        3.3.3.3            -          24960
*ai                   8.8.8.8            -          24960
*ai                   9.9.9.9            -          24961
*>i 7.7.7.7/32        7.7.7.7         24962(eth7)       24960
*>i 13.13.13.13/32    13.13.13.13     24960(eth7)       24960
*>i 14.14.14.14/32    3.3.3.3            -          24961
*ai                   8.8.8.8            -          24961
*ai                   9.9.9.9            -          24960
*>i 15.15.15.15/32    15.15.15.15     24961(eth7)       24960
BR7#
BR7#show ip bgp labeled-unicast all summary
BGP router identifier 10.12.28.169, local AS number 100
BGP table version is 7
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
3.3.3.3                  4   100  603        603       7      0      0  04:59:46               2
7.7.7.7                  4   100  602        607       7      0      0  04:59:43               1
8.8.8.8                  4   100  603        604       7      0      0  04:59:40               2
9.9.9.9                  4   100  603        604       7      0      0  04:59:47               2
13.13.13.13              4   100  602        608       7      0      0  04:59:42               1
15.15.15.15              4   100  602        606       7      0      0  04:59:40               1

Total number of neighbors 6

Total number of Established sessions 6
BR7#
BR7#show mpls ftn-table
 Primary FTN entry with FEC: 1.1.1.1/32, id: 25, row status: Active
  Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 20, in intf: - in label: 0 out-segment ix: 13
      Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 13, owner: BGP, Stale: NO, out intf: eth7, out label: 24960
    Nexthop addr: 3.3.3.3        cross connect ix: 20, op code: Push and Lookup


 Primary FTN entry with FEC: 3.3.3.3/32, id: 12, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth7, out label: 24320
    Nexthop addr: 29.1.1.159        cross connect ix: 3, op code: Push


 Primary FTN entry with FEC: 4.4.4.4/32, id: 13, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: eth7, out label: 3
    Nexthop addr: 29.1.1.159        cross connect ix: 4, op code: Push


 Primary FTN entry with FEC: 5.5.5.5/32, id: 9, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 15, in intf: - in label: 0 out-segment ix: 10
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 10, owner: LDP, Stale: NO, out intf: eth5, out label: 24332
    Nexthop addr: 33.1.1.161        cross connect ix: 15, op code: Push
```

```
Primary FTN entry with FEC: 6.6.6.6/32, id: 1, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth5, out label: 3
   Nexthop addr: 33.1.1.161        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 7.7.7.7/32, id: 10, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 17, in intf: - in label: 0 out-segment ix: 11
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 11, owner: LDP, Stale: NO, out intf: eth5, out label: 24336
   Nexthop addr: 33.1.1.161        cross connect ix: 17, op code: Push


Primary FTN entry with FEC: 7.7.7.7/32, id: 28, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 23, in intf: - in label: 0 out-segment ix: 16
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 16, owner: BGP, Stale: NO, out intf: eth5, out label: 24960
   Nexthop addr: 7.7.7.7        cross connect ix: 23, op code: Push


Primary FTN entry with FEC: 8.8.8.8/32, id: 14, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 8
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 8, owner: LDP, Stale: NO, out intf: eth7, out label: 24336
   Nexthop addr: 29.1.1.159        cross connect ix: 11, op code: Push


Primary FTN entry with FEC: 9.9.9.9/32, id: 24, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth7, out label: 24323
   Nexthop addr: 29.1.1.159        cross connect ix: 7, op code: Push


Primary FTN entry with FEC: 10.10.10.10/32, id: 16, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 7
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth7, out label: 24331
   Nexthop addr: 29.1.1.159        cross connect ix: 9, op code: Push


Primary FTN entry with FEC: 11.11.11.11/32, id: 17, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 9
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 9, owner: LDP, Stale: NO, out intf: eth7, out label: 24342
```

```
   Nexthop addr: 29.1.1.159        cross connect ix: 13, op code: Push


Primary FTN entry with FEC: 13.13.13.13/32, id: 11, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 19, in intf: - in label: 0 out-segment ix: 12
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 12, owner: LDP, Stale: NO, out intf: eth5, out label: 24342
   Nexthop addr: 33.1.1.161        cross connect ix: 19, op code: Push


Primary FTN entry with FEC: 13.13.13.13/32, id: 15, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 5, owner: BGP, Stale: NO, out intf: eth5, out label: 24960
   Nexthop addr: 13.13.13.13        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 14.14.14.14/32, id: 27, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 22, in intf: - in label: 0 out-segment ix: 15
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 15, owner: BGP, Stale: NO, out intf: eth7, out label: 24961
   Nexthop addr: 3.3.3.3        cross connect ix: 22, op code: Push and Lookup


Primary FTN entry with FEC: 15.15.15.15/32, id: 8, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth5, out label: 24322
   Nexthop addr: 33.1.1.161        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 15.15.15.15/32, id: 26, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 21, in intf: - in label: 0 out-segment ix: 14
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 14, owner: BGP, Stale: NO, out intf: eth5, out label: 24960
   Nexthop addr: 15.15.15.15        cross connect ix: 21, op code: Push


Primary FTN entry with FEC: 22.1.1.0/24, id: 18, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: eth7, out label: 3
   Nexthop addr: 29.1.1.159        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 24.1.1.0/24, id: 19, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
```

© 2026 IP Infusion Inc. Proprietary

```
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: eth7, out label: 3
   Nexthop addr: 29.1.1.159        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 25.1.1.0/24, id: 20, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: eth7, out label: 3
   Nexthop addr: 29.1.1.159        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 26.1.1.0/24, id: 2, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth5, out label: 3
   Nexthop addr: 33.1.1.161        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 27.1.1.0/24, id: 21, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: eth7, out label: 3
   Nexthop addr: 29.1.1.159        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 28.1.1.0/24, id: 3, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth5, out label: 3
   Nexthop addr: 33.1.1.161        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 30.1.1.0/24, id: 22, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: eth7, out label: 3
   Nexthop addr: 29.1.1.159        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 32.1.1.0/24, id: 4, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth5, out label: 3
   Nexthop addr: 33.1.1.161        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 40.1.1.0/24, id: 23, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
```

```
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: eth7, out label: 3
    Nexthop addr: 29.1.1.159        cross connect ix: 4, op code: Push


 Primary FTN entry with FEC: 42.1.1.0/24, id: 5, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth5, out label: 3
    Nexthop addr: 33.1.1.161        cross connect ix: 1, op code: Push


 Primary FTN entry with FEC: 50.1.1.0/24, id: 6, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth5, out label: 3
    Nexthop addr: 33.1.1.161        cross connect ix: 1, op code: Push


 Primary FTN entry with FEC: 60.1.1.0/24, id: 7, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth5, out label: 3
    Nexthop addr: 33.1.1.161        cross connect ix: 1, op code: Push


BR7#show ip bgp labeled-unicast 1.1.1.1/32
BGP routing table entry for 1.1.1.1/32
Paths: (3 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  7.7.7.7 15.15.15.15
  Local
    3.3.3.3 (metric 12) from 3.3.3.3 (10.12.28.158)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal, best
      Originator: 10.12.28.156, Cluster list: 10.12.28.158
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       13.13.13.13
      Last update: Mon Aug  6 00:10:31 2018

  Local
    8.8.8.8 (metric 12) from 8.8.8.8 (10.12.28.165)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal, add-path
      Originator: 10.12.28.156, Cluster list: 10.12.28.165
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       13.13.13.13
      Last update: Mon Aug  6 00:10:29 2018

  Local
    9.9.9.9 (metric 12) from 9.9.9.9 (10.12.28.166)
      Origin IGP, metric 0, localpref 100, label    24961, Out-label    24961, In-label  NA , valid,
internal, add-path
      Originator: 10.12.28.156, Cluster list: 10.12.28.166
      rx path_id: 0      tx path_id: 2
```

```
      Advertised to non peer-group peers:
       13.13.13.13
      Last update: Mon Aug  6 00:10:37 2018

BR7#
```

```
AGN4#show ip bgp labeled-unicast all

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.1/32       3.3.3.3            -           24960
* i                   3.3.3.3            -           24960
* i                   3.3.3.3            -           24960
* i                   3.3.3.3            -           24960
* i                   8.8.8.8            -           24960
* i                   8.8.8.8            -           24960
* i                   8.8.8.8            -           24960
* i                   8.8.8.8            -           24960
* i                   9.9.9.9            -           24961
* i                   9.9.9.9            -           24961
* i                   9.9.9.9            -           24961
*>i 7.7.7.7/32       7.7.7.7            -           24960
* i                   7.7.7.7            -           24960
* i                   7.7.7.7            -           24960
* i                   7.7.7.7            -           24960
*>  13.13.13.13/32   0.0.0.0         24960(eth6)          -
*>i 14.14.14.14/32    3.3.3.3            -           24961
* i                   3.3.3.3            -           24961
* i                   3.3.3.3            -           24961
* i                   3.3.3.3            -           24961
* i                   8.8.8.8            -           24961
* i                   8.8.8.8            -           24961
* i                   8.8.8.8            -           24961
* i                   9.9.9.9            -           24960
* i                   9.9.9.9            -           24960
* i                   9.9.9.9            -           24960
* i                   9.9.9.9            -           24960
*>i 15.15.15.15/32    15.15.15.15        -           24960
* i                   15.15.15.15        -           24960
* i                   15.15.15.15        -           24960
* i                   15.15.15.15        -           24960
AGN4#
AGN4#show ip bgp labeled-unicast all summary
BGP router identifier 10.12.28.170, local AS number 100
BGP table version is 7
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
5.5.5.5                  4   100  609        603       7      0      0  05:00:20               8
10.10.10.10              4   100  608        603       7      0      0  05:00:22               8
11.11.11.11              4   100  607        602       7      0      0  05:00:00               6
12.12.12.12              4   100  609        603       7      0      0  05:00:29               8

Total number of neighbors 4

Total number of Established sessions 4
AGN4#
AGN4#show ip bgp labeled-unicast 1.1.1.1/32
BGP routing table entry for 1.1.1.1/32
Paths: (11 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  Local
    3.3.3.3 (metric 20) from 5.5.5.5 (10.12.28.160)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal, best
      Originator: 10.12.28.158, Cluster list: 10.12.28.160 10.12.28.158
      rx path_id: 1      tx path_id: -1
```

```
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:36 2018

  Local
    3.3.3.3 (metric 20) from 10.10.10.10 (10.12.28.167)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal
      Originator: 10.12.28.158, Cluster list: 10.12.28.167 10.12.28.158
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:34 2018

  Local
    3.3.3.3 (metric 20) from 11.11.11.11 (10.12.28.168)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal
      Originator: 10.12.28.158, Cluster list: 10.12.28.168 10.12.28.158
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:56 2018

  Local
    3.3.3.3 (metric 20) from 12.12.12.12 (10.12.28.169)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal
      Originator: 10.12.28.158, Cluster list: 10.12.28.169 10.12.28.158
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:31 2018

  Local
    8.8.8.8 (metric 20) from 5.5.5.5 (10.12.28.160)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal
      Originator: 10.12.28.165, Cluster list: 10.12.28.160 10.12.28.165
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:40 2018

  Local
    8.8.8.8 (metric 20) from 10.10.10.10 (10.12.28.167)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal
      Originator: 10.12.28.165, Cluster list: 10.12.28.167 10.12.28.165
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:38 2018

  Local
    8.8.8.8 (metric 20) from 11.11.11.11 (10.12.28.168)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal
      Originator: 10.12.28.165, Cluster list: 10.12.28.168 10.12.28.165
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:11:08 2018

  Local
    8.8.8.8 (metric 20) from 12.12.12.12 (10.12.28.169)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal
      Originator: 10.12.28.165, Cluster list: 10.12.28.169 10.12.28.165
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:31 2018

  Local
    9.9.9.9 (metric 20) from 5.5.5.5 (10.12.28.160)
      Origin IGP, metric 0, localpref 100, label    24961, Out-label    24961, In-label  NA , valid,
```

```
internal
      Originator: 10.12.28.166, Cluster list: 10.12.28.160 10.12.28.166
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:40 2018

  Local
    9.9.9.9 (metric 20) from 10.10.10.10 (10.12.28.167)
      Origin IGP, metric 0, localpref 100, label    24961, Out-label    24961, In-label  NA , valid,
internal
      Originator: 10.12.28.166, Cluster list: 10.12.28.167 10.12.28.166
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:54 2018

  Local
    9.9.9.9 (metric 20) from 12.12.12.12 (10.12.28.169)
      Origin IGP, metric 0, localpref 100, label    24961, Out-label    24961, In-label  NA , valid,
internal
      Originator: 10.12.28.166, Cluster list: 10.12.28.169 10.12.28.166
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:43 2018

AGN4#show mpls ftn-table
 Primary FTN entry with FEC: 1.1.1.1/32, id: 19, row status: Active
  Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 14, in intf: - in label: 0 out-segment ix: 13
      Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 13, owner: BGP, Stale: NO, out intf: eth6, out label: 24960
    Nexthop addr: 3.3.3.3        cross connect ix: 14, op code: Push and Lookup


 Primary FTN entry with FEC: 3.3.3.3/32, id: 1, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth6, out label: 24325
    Nexthop addr: 32.1.1.161        cross connect ix: 2, op code: Push


 Primary FTN entry with FEC: 5.5.5.5/32, id: 2, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 2
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth6, out label: 24331
    Nexthop addr: 32.1.1.161        cross connect ix: 3, op code: Push


 Primary FTN entry with FEC: 6.6.6.6/32, id: 3, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth6, out label: 3
    Nexthop addr: 32.1.1.161        cross connect ix: 4, op code: Push


 Primary FTN entry with FEC: 7.7.7.7/32, id: 4, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
```

```
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 4
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: eth6, out label: 24334
    Nexthop addr: 32.1.1.161        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 7.7.7.7/32, id: 20, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 15, in intf: - in label: 0 out-segment ix: 14
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 14, owner: BGP, Stale: NO, out intf: eth6, out label: 24960
   Nexthop addr: 7.7.7.7        cross connect ix: 15, op code: Push


Primary FTN entry with FEC: 8.8.8.8/32, id: 5, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth6, out label: 24327
   Nexthop addr: 32.1.1.161        cross connect ix: 6, op code: Push


Primary FTN entry with FEC: 9.9.9.9/32, id: 6, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth6, out label: 24329
   Nexthop addr: 32.1.1.161        cross connect ix: 7, op code: Push


Primary FTN entry with FEC: 10.10.10.10/32, id: 7, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 7
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth6, out label: 24323
   Nexthop addr: 32.1.1.161        cross connect ix: 8, op code: Push


Primary FTN entry with FEC: 11.11.11.11/32, id: 8, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 10
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 10, owner: LDP, Stale: NO, out intf: eth6, out label: 24348
   Nexthop addr: 32.1.1.161        cross connect ix: 11, op code: Push


Primary FTN entry with FEC: 12.12.12.12/32, id: 9, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 8
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 8, owner: LDP, Stale: NO, out intf: eth6, out label: 24320
   Nexthop addr: 32.1.1.161        cross connect ix: 9, op code: Push


Primary FTN entry with FEC: 14.14.14.14/32, id: 21, row status: Active
```

```
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 16, in intf: - in label: 0 out-segment ix: 15
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 15, owner: BGP, Stale: NO, out intf: eth6, out label: 24961
   Nexthop addr: 3.3.3.3        cross connect ix: 16, op code: Push and Lookup


Primary FTN entry with FEC: 15.15.15.15/32, id: 10, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 9
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 9, owner: LDP, Stale: NO, out intf: eth6, out label: 24321
   Nexthop addr: 32.1.1.161        cross connect ix: 10, op code: Push


Primary FTN entry with FEC: 15.15.15.15/32, id: 18, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 12
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 12, owner: BGP, Stale: NO, out intf: eth6, out label: 24960
   Nexthop addr: 15.15.15.15        cross connect ix: 13, op code: Push


Primary FTN entry with FEC: 26.1.1.0/24, id: 11, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth6, out label: 3
   Nexthop addr: 32.1.1.161        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 28.1.1.0/24, id: 12, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth6, out label: 3
   Nexthop addr: 32.1.1.161        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 33.1.1.0/24, id: 13, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth6, out label: 3
   Nexthop addr: 32.1.1.161        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 42.1.1.0/24, id: 14, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth6, out label: 3
   Nexthop addr: 32.1.1.161        cross connect ix: 4, op code: Push
```

```
 Primary FTN entry with FEC: 50.1.1.0/24, id: 15, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth6, out label: 3
    Nexthop addr: 32.1.1.161        cross connect ix: 4, op code: Push


 Primary FTN entry with FEC: 60.1.1.0/24, id: 16, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth6, out label: 3
    Nexthop addr: 32.1.1.161        cross connect ix: 4, op code: Push


AGN4#
```

```
AGN2#show ip bgp labeled-unicast all

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.1/32        3.3.3.3            -          24960
* i                   3.3.3.3            -          24960
* i                   3.3.3.3            -          24960
* i                   3.3.3.3            -          24960
* i                   8.8.8.8            -          24960
* i                   8.8.8.8            -          24960
* i                   8.8.8.8            -          24960
* i                   9.9.9.9            -          24961
* i                   9.9.9.9            -          24961
*>  7.7.7.7/32       0.0.0.0         24960(eth2)          -
*>i 13.13.13.13/32    13.13.13.13        -          24960
* i                   13.13.13.13        -          24960
* i                   13.13.13.13        -          24960
* i                   13.13.13.13        -          24960
*>i 14.14.14.14/32    3.3.3.3            -          24961
* i                   3.3.3.3            -          24961
* i                   3.3.3.3            -          24961
* i                   3.3.3.3            -          24961
* i                   8.8.8.8            -          24961
* i                   8.8.8.8            -          24961
* i                   9.9.9.9            -          24960
* i                   9.9.9.9            -          24960
* i                   9.9.9.9            -          24960
*>i 15.15.15.15/32    15.15.15.15        -          24960
* i                   15.15.15.15        -          24960
* i                   15.15.15.15        -          24960
* i                   15.15.15.15        -          24960
AGN2#show ip bgp labeled-unicast all summary
BGP router identifier 10.12.28.162, local AS number 100
BGP table version is 8
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
5.5.5.5                  4   100  579        573       8      0      0  04:45:08               8
10.10.10.10              4   100  578        573       8      0      0  04:45:07               8
11.11.11.11              4   100  578        573       8      0      0  04:45:05               6
12.12.12.12              4   100  578        573       8      0      0  04:45:07               4

Total number of neighbors 4
```

```
Total number of Established sessions 4
AGN2#
AGN2#show ip bgp labeled-unicast 1.1.1.1/32
BGP routing table entry for 1.1.1.1/32
Paths: (9 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  Local
    3.3.3.3 (metric 20) from 5.5.5.5 (10.12.28.160)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal, best
      Originator: 10.12.28.158, Cluster list: 10.12.28.160 10.12.28.158
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:32 2018

  Local
    3.3.3.3 (metric 20) from 10.10.10.10 (10.12.28.167)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal
      Originator: 10.12.28.158, Cluster list: 10.12.28.167 10.12.28.158
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:34 2018

  Local
    3.3.3.3 (metric 20) from 11.11.11.11 (10.12.28.168)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal
      Originator: 10.12.28.158, Cluster list: 10.12.28.168 10.12.28.158
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:32 2018

  Local
    3.3.3.3 (metric 20) from 12.12.12.12 (10.12.28.169)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal
      Originator: 10.12.28.158, Cluster list: 10.12.28.169 10.12.28.158
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:34 2018

  Local
    8.8.8.8 (metric 20) from 5.5.5.5 (10.12.28.160)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal, backup
      Originator: 10.12.28.165, Cluster list: 10.12.28.160 10.12.28.165
      rx path_id: 0      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:28 2018

  Local
    8.8.8.8 (metric 20) from 10.10.10.10 (10.12.28.167)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal
      Originator: 10.12.28.165, Cluster list: 10.12.28.167 10.12.28.165
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:37 2018

  Local
    8.8.8.8 (metric 20) from 11.11.11.11 (10.12.28.168)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal
      Originator: 10.12.28.165, Cluster list: 10.12.28.168 10.12.28.165
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:11:07 2018
```

```
  Local
    9.9.9.9 (metric 20) from 5.5.5.5 (10.12.28.160)
      Origin IGP, metric 0, localpref 100, label    24961, Out-label    24961, In-label  NA , valid,
internal
      Originator: 10.12.28.166, Cluster list: 10.12.28.160 10.12.28.166
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:39 2018

  Local
    9.9.9.9 (metric 20) from 10.10.10.10 (10.12.28.167)
      Origin IGP, metric 0, localpref 100, label    24961, Out-label    24961, In-label  NA , valid,
internal
      Originator: 10.12.28.166, Cluster list: 10.12.28.167 10.12.28.166
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:52 2018


BGP PIC Enabled
AGN2#

AGN2#sh mpls ftn-table
 Primary FTN entry with FEC: 1.1.1.1/32, id: 19, row status: Active
  Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 12
      Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 12, owner: BGP, Stale: NO, out intf: eth2, out label: 24960
    Nexthop addr: 3.3.3.3        cross connect ix: 13, op code: Push and Lookup


 Non-primary FTN entry with FEC: 1.1.1.1/32, id: 24, row status: Active
  Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 17, in intf: - in label: 0 out-segment ix: 16
      Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 16, owner: BGP, Stale: NO, out intf: eth2, out label: 24960
    Nexthop addr: 8.8.8.8        cross connect ix: 17, op code: Push and Lookup


 Primary FTN entry with FEC: 3.3.3.3/32, id: 1, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth2, out label: 24325
    Nexthop addr: 60.1.1.161        cross connect ix: 2, op code: Push


 Primary FTN entry with FEC: 5.5.5.5/32, id: 2, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 2
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth2, out label: 24331
    Nexthop addr: 60.1.1.161        cross connect ix: 3, op code: Push


 Primary FTN entry with FEC: 6.6.6.6/32, id: 3, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
```

```
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth2, out label: 3
   Nexthop addr: 60.1.1.161        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 8.8.8.8/32, id: 4, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 4
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: eth2, out label: 24327
   Nexthop addr: 60.1.1.161        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 9.9.9.9/32, id: 5, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth2, out label: 24329
   Nexthop addr: 60.1.1.161        cross connect ix: 6, op code: Push


Primary FTN entry with FEC: 10.10.10.10/32, id: 6, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth2, out label: 24323
   Nexthop addr: 60.1.1.161        cross connect ix: 7, op code: Push


Primary FTN entry with FEC: 11.11.11.11/32, id: 7, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 10
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 10, owner: LDP, Stale: NO, out intf: eth2, out label: 24345
   Nexthop addr: 60.1.1.161        cross connect ix: 11, op code: Push


Primary FTN entry with FEC: 12.12.12.12/32, id: 8, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 7
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth2, out label: 24320
   Nexthop addr: 60.1.1.161        cross connect ix: 8, op code: Push


Primary FTN entry with FEC: 13.13.13.13/32, id: 9, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 9
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 9, owner: LDP, Stale: NO, out intf: eth2, out label: 24339
   Nexthop addr: 60.1.1.161        cross connect ix: 10, op code: Push


Primary FTN entry with FEC: 13.13.13.13/32, id: 17, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
```

```
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 12, in intf: - in label: 0 out-segment ix: 11
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 11, owner: BGP, Stale: NO, out intf: eth2, out label: 24960
   Nexthop addr: 13.13.13.13        cross connect ix: 12, op code: Push


Primary FTN entry with FEC: 14.14.14.14/32, id: 22, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 16, in intf: - in label: 0 out-segment ix: 15
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 15, owner: BGP, Stale: NO, out intf: eth2, out label: 24961
   Nexthop addr: 3.3.3.3        cross connect ix: 16, op code: Push and Lookup


Non-primary FTN entry with FEC: 14.14.14.14/32, id: 25, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 18, in intf: - in label: 0 out-segment ix: 17
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 17, owner: BGP, Stale: NO, out intf: eth2, out label: 24961
   Nexthop addr: 8.8.8.8        cross connect ix: 18, op code: Push and Lookup


Primary FTN entry with FEC: 15.15.15.15/32, id: 10, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 8
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 8, owner: LDP, Stale: NO, out intf: eth2, out label: 24321
   Nexthop addr: 60.1.1.161        cross connect ix: 9, op code: Push


Primary FTN entry with FEC: 15.15.15.15/32, id: 26, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 19, in intf: - in label: 0 out-segment ix: 18
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 18, owner: BGP, Stale: NO, out intf: eth2, out label: 24960
   Nexthop addr: 15.15.15.15        cross connect ix: 19, op code: Push


Primary FTN entry with FEC: 26.1.1.0/24, id: 11, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth2, out label: 3
   Nexthop addr: 60.1.1.161        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 28.1.1.0/24, id: 12, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth2, out label: 3
   Nexthop addr: 60.1.1.161        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 32.1.1.0/24, id: 13, row status: Active
```

```
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth2, out label: 3
    Nexthop addr: 60.1.1.161        cross connect ix: 4, op code: Push


 Primary FTN entry with FEC: 33.1.1.0/24, id: 14, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth2, out label: 3
    Nexthop addr: 60.1.1.161        cross connect ix: 4, op code: Push


 Primary FTN entry with FEC: 42.1.1.0/24, id: 15, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth2, out label: 3
    Nexthop addr: 60.1.1.161        cross connect ix: 4, op code: Push


 Primary FTN entry with FEC: 50.1.1.0/24, id: 16, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth2, out label: 3
    Nexthop addr: 60.1.1.161        cross connect ix: 4, op code: Push


AGN2#
```

```
AGN5#show ip bgp labeled-unicast all

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, S - stale
     Network          Next Hop        In Label     Out Label
*>i 1.1.1.1/32        3.3.3.3            -           24960
* i                   3.3.3.3            -           24960
* i                   3.3.3.3            -           24960
* i                   3.3.3.3            -           24960
*>i 7.7.7.7/32        7.7.7.7            -           24960
* i                   7.7.7.7            -           24960
* i                   7.7.7.7            -           24960
* i                   7.7.7.7            -           24960
*>i 13.13.13.13/32    13.13.13.13        -           24960
* i                   13.13.13.13        -           24960
* i                   13.13.13.13        -           24960
* i                   13.13.13.13        -           24960
*>i 14.14.14.14/32    3.3.3.3            -           24961
* i                   8.8.8.8            -           24961
* i                   9.9.9.9            -           24960
* i                   9.9.9.9            -           24960
*>  15.15.15.15/32   0.0.0.0         24960(eth7)          -
AGN5#
AGN5#show ip bgp labeled-unicast all summary
BGP router identifier 10.12.28.172, local AS number 100
BGP table version is 8
1 BGP AS-PATH entries
0 BGP community entries
```

```
Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
5.5.5.5                  4   100  610        606       8      0      0  05:01:41               4
10.10.10.10              4   100  609        606       8      0      0  05:01:54               4
11.11.11.11              4   100  609        606       8      0      0  05:01:39               4
12.12.12.12              4   100  610        606       8      0      0  05:01:39               4

Total number of neighbors 4

Total number of Established sessions 4
AGN5#
AGN5#show ip bgp labeled-unicast 1.1.1.1/32
BGP routing table entry for 1.1.1.1/32
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  Local
    3.3.3.3 (metric 20) from 5.5.5.5 (10.12.28.160)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal, best
      Originator: 10.12.28.158, Cluster list: 10.12.28.160 10.12.28.158
      rx path_id: -1      tx path_id: -1
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:34 2018

  Local
    3.3.3.3 (metric 20) from 10.10.10.10 (10.12.28.167)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal
      Originator: 10.12.28.158, Cluster list: 10.12.28.167 10.12.28.158
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:32 2018

  Local
    3.3.3.3 (metric 20) from 11.11.11.11 (10.12.28.168)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal
      Originator: 10.12.28.158, Cluster list: 10.12.28.168 10.12.28.158
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:33 2018

  Local
    3.3.3.3 (metric 20) from 12.12.12.12 (10.12.28.169)
      Origin IGP, metric 0, localpref 100, label    24960, Out-label    24960, In-label  NA , valid,
internal
      Originator: 10.12.28.158, Cluster list: 10.12.28.169 10.12.28.158
      Not advertised to any peer
      Last update: Mon Aug  6 00:10:33 2018


BGP PIC Enabled
AGN5#show mpls ftn-table
 Primary FTN entry with FEC: 1.1.1.1/32, id: 19, row status: Active
  Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 14, in intf: - in label: 0 out-segment ix: 13
      Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 13, owner: BGP, Stale: NO, out intf: eth7, out label: 24960
    Nexthop addr: 3.3.3.3        cross connect ix: 14, op code: Push and Lookup


 Primary FTN entry with FEC: 3.3.3.3/32, id: 11, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: eth7, out label: 24326
```

© 2026 IP Infusion Inc. Proprietary

```
   Nexthop addr: 42.1.1.161        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 5.5.5.5/32, id: 9, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 7
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: eth7, out label: 24333
   Nexthop addr: 42.1.1.161        cross connect ix: 8, op code: Push


Primary FTN entry with FEC: 6.6.6.6/32, id: 1, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth7, out label: 3
   Nexthop addr: 42.1.1.161        cross connect ix: 1, op code: Push


Primary FTN entry with FEC: 7.7.7.7/32, id: 14, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 8
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 8, owner: LDP, Stale: NO, out intf: eth7, out label: 24337
   Nexthop addr: 42.1.1.161        cross connect ix: 9, op code: Push


Primary FTN entry with FEC: 7.7.7.7/32, id: 21, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 15, in intf: - in label: 0 out-segment ix: 14
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 14, owner: BGP, Stale: NO, out intf: eth7, out label: 24960
   Nexthop addr: 7.7.7.7        cross connect ix: 15, op code: Push


Primary FTN entry with FEC: 8.8.8.8/32, id: 12, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: eth7, out label: 24328
   Nexthop addr: 42.1.1.161        cross connect ix: 5, op code: Push


Primary FTN entry with FEC: 9.9.9.9/32, id: 13, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 6
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: eth7, out label: 24330
   Nexthop addr: 42.1.1.161        cross connect ix: 6, op code: Push


Primary FTN entry with FEC: 10.10.10.10/32, id: 10, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3
```

```
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: eth7, out label: 24324
   Nexthop addr: 42.1.1.161        cross connect ix: 3, op code: Push


Primary FTN entry with FEC: 11.11.11.11/32, id: 15, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 10
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 10, owner: LDP, Stale: NO, out intf: eth7, out label: 24349
   Nexthop addr: 42.1.1.161        cross connect ix: 11, op code: Push


Primary FTN entry with FEC: 12.12.12.12/32, id: 2, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: eth7, out label: 24320
   Nexthop addr: 42.1.1.161        cross connect ix: 2, op code: Push


Primary FTN entry with FEC: 13.13.13.13/32, id: 16, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 9
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 9, owner: LDP, Stale: NO, out intf: eth7, out label: 24343
   Nexthop addr: 42.1.1.161        cross connect ix: 10, op code: Push


Primary FTN entry with FEC: 13.13.13.13/32, id: 17, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 12, in intf: - in label: 0 out-segment ix: 11
     Owner: BGP, Persistent: No, Admin Status: Down, Oper Status: Down
      Out-segment with ix: 11, owner: BGP, Stale: NO, out intf: eth7, out label: 24960
   Nexthop addr: 13.13.13.13        cross connect ix: 12, op code: Push


Primary FTN entry with FEC: 14.14.14.14/32, id: 20, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 16, in intf: - in label: 0 out-segment ix: 15
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 15, owner: BGP, Stale: NO, out intf: eth7, out label: 24961
   Nexthop addr: 3.3.3.3        cross connect ix: 16, op code: Push and Lookup


Non-primary FTN entry with FEC: 14.14.14.14/32, id: 22, row status: Active
 Owner: BGP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
   Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Primary:  Cross connect ix: 17, in intf: - in label: 0 out-segment ix: 16
     Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 16, owner: BGP, Stale: NO, out intf: eth7, out label: 24961
   Nexthop addr: 8.8.8.8        cross connect ix: 17, op code: Push and Lookup


Primary FTN entry with FEC: 26.1.1.0/24, id: 3, row status: Active
 Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0,   Protected LSP id: 0, Description: N/A
```

```
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth7, out label: 3
    Nexthop addr: 42.1.1.161        cross connect ix: 1, op code: Push


 Primary FTN entry with FEC: 28.1.1.0/24, id: 4, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth7, out label: 3
    Nexthop addr: 42.1.1.161        cross connect ix: 1, op code: Push


 Primary FTN entry with FEC: 32.1.1.0/24, id: 5, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth7, out label: 3
    Nexthop addr: 42.1.1.161        cross connect ix: 1, op code: Push


 Primary FTN entry with FEC: 33.1.1.0/24, id: 6, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth7, out label: 3
    Nexthop addr: 42.1.1.161        cross connect ix: 1, op code: Push


 Primary FTN entry with FEC: 50.1.1.0/24, id: 7, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth7, out label: 3
    Nexthop addr: 42.1.1.161        cross connect ix: 1, op code: Push


 Primary FTN entry with FEC: 60.1.1.0/24, id: 8, row status: Active
  Owner: LDP, Stale: NO, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0,   Protected LSP id: 0, Description: N/A
    Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
  Primary:  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: eth7, out label: 3
    Nexthop addr: 42.1.1.161        cross connect ix: 1, op code: Push


AGN5#
```

BGP4+ Additional Paths Configuration

Overview

The Border Gateway Protocol (BGP) ADDPATH feature allows the advertisement of multiple paths through the same peering 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.

Normal BGP Behavior

By default, all BGP routers and Route-Reflectors propagate only their best paths over their sessions. In case they advertise any route with the same NLRI as a previously advertised route, the latest one implicitly replaces the previous advertisement, which is known as an Implicit Withdraw. The Implicit Withdraw can achieve better scaling, but at the cost of path diversity.

The use of route-reflectors (or confederations), thus has significant effect on redundancy by hiding alternate paths. Using full-mesh is not an option, so a mechanism is needed to allow the propagation of multiple alternate paths in an RR/Confederation environment. Such mechanism is already available in BGP/MPLS VPN scenarios, where multiple point of attachments for CE sites could utilize different RD values to differentiate the same routes advertised from different connection points. However, a generic solution is required, allowing for advertising multiple alternate paths with IPv4 or any other address-family.

The “Advertisement of Multiple Paths in BGP” or “BGP Add-Path” as the feature is usually called is a BGP extension that allows the advertisement of multiple paths for the same address prefix without the new paths implicitly replacing any previously advertised ones.

BGP Behavior with Additional Paths

The advertisement of multiple paths in BGP is made possible by aending a BGP OPEN message to the neighbor with a BGP capability code of 69, which identifies the BGP ADD-PATH Capability.

Address Family Identifier(AFI)	2 octets
Subsequent Address Family Identifier(SAFI)	1 octet
Send/Receive	1 octet

The send/receive field in the BGP Capability TLV indicates whether for a given <AFI, SAFI>, the sender is able to:

- Receive multiple paths from its peer (value 1)
- Send multiple paths to its peer (value 2), or
- both (value 3)
- Each alternate path is identified by a Path Identifier in addition to the address prefix

Path Identifier	4 octets
Length	1 octet
Prefix	variable

In the event of a next-hop failure, the BGP Add-Path feature hence improves the BGP control plane convergence

© 2026 IP Infusion Inc. Proprietary

Note: BGP does not advertise additional paths (non-best paths) to a neighbor when the neighbor is an eBGP peer with implicit next-hop-self or when next-hop-self is explicitly configured using the neighbor A.B.C.D next-hop-self command.

© 2026 IP Infusion Inc. Proprietary

Topology

Figure 87. BGP4+ additional paths

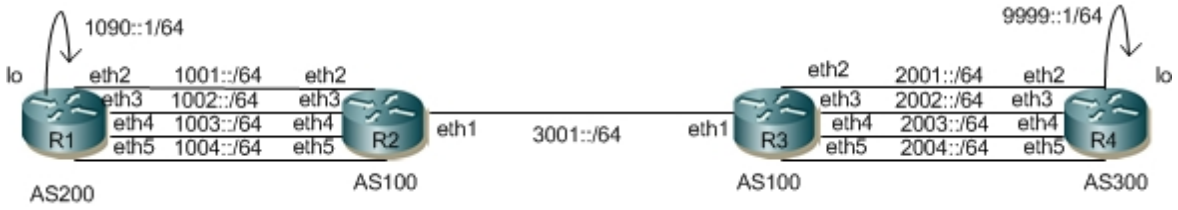

R1

`#configure terminal`	Enter the configure mode.
`(config)#interface eth2`	Enter interface mode for interface eth2
`(config-if)#ipv6 address 1001::1/64`	Configure IPv6 address for the interface eth2
`(config-if)#exit`	Exit the interface mode
`(config)#interface eth3`	Enter interface mode for interface eth3
`(config-if)#ipv6 address 1002::1/64`	Configure IPv6 address for the interface eth3
`(config-if)#exit`	Exit the interface mode
`(config)#interface eth4`	Enter interface mode for interface eth4
`(config-if)#ipv6 address 1003::1/64`	Configure IPv6 address for the interface eth4
`(config-if)#exit`	Exit the interface mode
`(config)#interface eth5`	Enter interface mode for interface eth5
`(config-if)#ipv6 address 1004::1/64`	Configure IPv6 address for the interface eth5
`(config-if)#exit`	Exit the interface mode
`(config)#interface lo`	Enter interface mode for loopback lo
`(config-if)#ipv6 address 1090::1/64`	Configure IPv6 address for Loopback interface lo
`(config-if)#exit`	Exit the interface mode
`(config)#router bgp 200`	Enter the router bgp mode
`(config-router)#neighbor 1001::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`(config-router)#neighbor 1002::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`(config-router)#neighbor 1004::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`(config-router)#neighbor 1003::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.

© 2026 IP Infusion Inc. Proprietary

(config-router-af)#neighbor 1001::2 activate	Activate the neighbor router with peer address.
(config-router-af)#neighbor 1002::2 activate	Activate the neighbor router with peer address.
(config-router-af)#neighbor 1003::2 activate	Activate the neighbor router with peer address.
(config-router-af)#neighbor 1004::2 activate	Activate the neighbor router with peer address.
(config-router-af)#network 1090::/64	Activate the neighbor router with peer address.
(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 commit
(config)#exit	Exit the configure mode

R2

#configure terminal	Enter the cnfigure mode.
(config)#interface eth1	Enter interface mode for interface eth1
(config-if)#ipv6 address 3001::1/64	Configure IPv6 address for the interface eth1
(config-if)#ipv6 router ospf area 0	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
(config-if)#exit	Exit the interface mode
(config)#interface eth2	Enter interface mode for interface eth2
(config-if)#ipv6 address 1001::2/64	Configure IPv6 address for the interface eth2
(config-if)#exit	Exit the interface mode
(config)#interface eth3	Enter interface mode for interface eth3
(config-if)#ipv6 address 1002::2/64	Configure IPv6 address for the interface eth3
(config-if)#exit	Exit the interface mode
(config)#interface eth4	Enter interface mode for interface eth4
(config-if)#ipv6 address 1003::2/64	Configure IPv6 address for the interface eth4
(config-if)#exit	Exit the interface mode
(config)#interface eth5	Enter interface mode for interface eth5
(config-if)#ipv6 address 1004::2/64	Configure IPv6 address for the interface eth5
(config-if)#exit	Exit the interface mode
(config)#router bgp 100	Enter the router bgp mode
(config-router)#neighbor 3001::2 remote-as 100	Specify a neighbor router with peer address and remote-as for BGP peering.
(config-router)#neighbor 1001::1 remote-as 200	Specify a neighbor router with peer address and remote-as for BGP peering.
(config-router)#neighbor 1002::1 remote-as 200	Specify a neighbor router with peer address and remote-as for BGP peering.

© 2026 IP Infusion Inc. Proprietary

(config-router)#neighbor 1003::1 remote-as 200	Specify a neighbor router with peer address and remote-as for BGP peering.
(config-router)#neighbor 1004::1 remote-as 200	Specify a neighbor router with peer address and remote-as for BGP peering.
(config-router)#address-family ipv6 unicast	Enter address-family mode for neighbor router session to activate.
(config-router-af)#neighbor 1001::1 activate	Activate the neighbor router with peer address.
(config-router-af)#neighbor 1002::1 activate	Activate the neighbor router with peer address.
(config-router-af)#neighbor 1003::1 activate	Activate the neighbor router with peer address.
(config-router-af)#neighbor 1004::1 activate	Activate the neighbor router with peer address.
(config-router-af)#neighbor 3001::2 activate	Activate the neighbor router with peer address.
(config-router-af)#exit-address-family	Exit address family mode.
(config-router)#exit	Exit the router BGP mode and enter the config mode
(config)#router ipv6 ospf	Enter Router OSPFv3 mode.
(config-router)#redistribute connected	Configure Redistribution of Connected networks into OSPF
(config-router)#exit	Exit the router ospf mode and enter the config mode
(config)#commit	Apply commit
(config)#exit	Exit the configure mode

R3

#configure terminal	Enter the configure mode.
(config)#interface eth1	Enter interface mode for interface eth1
(config-if)#ipv6 address 3001::2/64	Configure IPv6 address for the interface eth1
(config-if)#ipv6 router ospf area 0	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
(config-if)#exit	Exit the interface mode
(config)#interface eth2	Enter interface mode for interface eth2
(config-if)#ipv6 address 2001::2/64	Configure IPv6 address for the interface eth2
(config-if)#exit	Exit the interface mode
(config)#interface eth3	Enter interface mode for interface eth3
(config-if)#ipv6 address 2002::2/64	Configure IPv6 address for the interface eth3
(config-if)#exit	Exit the interface mode
(config)#interface eth4	Enter interface mode for interface eth4
(config-if)#ipv6 address 2003::2/64	Configure IPv6 address for the interface eth4
(config-if)#exit	Exit the interface mode
(config)#interface eth5	Enter interface mode for interface eth5

© 2026 IP Infusion Inc. Proprietary

(config-if)#ipv6 address 2004::2/64	Configure IPv6 address for the interface eth5
(config-if)#exit	Exit the interface mode
(config)#router bgp 100	Enter the router bgp mode
(config-router)#neighbor 3001::1 remote-as 100	Specify a neighbor router with peer address and remote-as for BGP peering.
(config-router)#neighbor 2001::1 remote-as 300	Specify a neighbor router with peer address and remote-as for BGP peering.
(config-router)#neighbor 2002::1 remote-as 300	Specify a neighbor router with peer address and remote-as for BGP peering.
(config-router)#neighbor 2003::1 remote-as 300	Specify a neighbor router with peer address and remote-as for BGP peering.
(config-router)#neighbor 2004::1 remote-as 300	Specify a neighbor router with peer address and remote-as for BGP peering.
(config-router)#address-family ipv6 unicast	Enter address-family mode for neighbor router session to activate.
(config-router-af)#neighbor 2001::1 activate	Activate the neighbor router with peer address.
(config-router-af)#neighbor 2002::1 activate	Activate the neighbor router with peer address.
(config-router-af)#neighbor 2003::1 activate	Activate the neighbor router with peer address.
(config-router-af)#neighbor 2004::1 activate	Activate the neighbor router with peer address.
(config-router-af)#neighbor 3001::1 activate	Activate the neighbor router with peer address.
(config-router-af)#exit-address-family	Exit address family mode.
(config-router)#exit	Exit Router BGP mode
(config)#router ipv6 ospf	Enter Router OSPFv3 mode.
(config-router)#redistribute connected	Configure Redistribution of Connected networks into OSPF
(config-router)#exit	Exit the router ospf mode and enter the configure mode
(config)#commit	Apply commit
(config)#exit	Exit the configure mode

R4

#configure terminal	Enter the configure mode.
(config)#interface eth2	Enter interface mode for interface eth2
(config-if)#ipv6 address 2001::1/64	Configure IPv6 address for the interface eth2
(config-if)#exit	Exit the interface mode
(config)#interface eth3	Enter interface mode for interface eth3
(config-if)#ipv6 address 2002::1/64	Configure IPv6 address for the interface eth3
(config-if)#exit	Exit the interface mode

© 2026 IP Infusion Inc. Proprietary

`(config)#interface eth4`	Enter interface mode for interface eth4
`(config-if)#ipv6 address 2003::1/64`	Configure IPv6 address for the interface eth4
`(config-if)#exit`	Exit the interface mode
`(config)#interface eth5`	Enter interface mode for interface eth5
`(config-if)#ipv6 address 2004::1/64`	Configure IPv6 address for the interface eth5
`(config-if)#exit`	Exit the interface mode
`(config)#interface lo`	Enter interface mode for loopback lo
`(config-if)#ipv6 address 9999::1/64`	Configure IPv6 address for Loopback interface lo
`(config-if)#exit`	Exit the interface mode
`(config)#router bgp 300`	Enter the router bgp mode
`(config-router)#neighbor 2001::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`(config-router)#neighbor 2002::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`(config-router)#neighbor 2003::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`(config-router)#neighbor 2004::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`(config-router-af)#neighbor 2001::2 activate`	Activate the neighbor router with peer address.
`(config-router-af)#neighbor 2002::2 activate`	Activate the neighbor router with peer address.
`(config-router-af)#neighbor 2003::2 activate`	Activate the neighbor router with peer address.
`(config-router-af)#neighbor 2004::2 activate`	Activate the neighbor router with peer address.
`(config-router-af)#network 9999::/64`	Activate the neighbor router with peer address.
`(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 commit
`(config)#exit`	Exit the configure mode

© 2026 IP Infusion Inc. Proprietary

Additional Paths at the Global Level

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`	Configure R2 to send additional paths to 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 commit
`(config)#exit`	Exit the configure mode

R3

`#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 receive`	Configure R3 to receive additional paths from 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 commit
`(config)#exit`	Exit the configure mode

© 2026 IP Infusion Inc. Proprietary

Additional Paths at the Neighbor Level

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)#neighbor 3001::2 additional-paths send`	Configure R2 to send additional paths to all iBGP neighbors
`(config-router-af)#neighbor 3001::2 advertise additional-paths 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 commit
`(config)#exit`	Exit the configure mode

R3

`#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)#neighbor 3001::1 additional-paths receive`	Configure R3 to receive additional paths from 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 commit
`(config)#exit`	Exit the configure mode

© 2026 IP Infusion Inc. Proprietary

Validation

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 00:14:55
  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

```
#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
```

© 2026 IP Infusion Inc. Proprietary

Additional Paths Send and Receive at Address-family level

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 commit
`(config)#exit`	Exit the configure mode

R3

`#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 R3 to send additional paths to and receive additional paths from all iBGP neighbors
`(config-router-af)#bgp additional-paths select all`	Configure R3 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 commit
`(config)#exit`	Exit the configure mode

© 2026 IP Infusion Inc. Proprietary

Additional Paths at the Neighbor Level

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)#neighbor 3001::2 additional-paths send-receive`	Configure R2 to send-receive additional paths to the iBGP neighbor R3
`(config-router-af)#neighbor 3001::2 advertise additional-paths all`	Configure R2 to advertise all available paths to the iBGP neighbor R3
`(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 commit
`(config)#exit`	Exit the configure mode

R3

`#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)#neighbor 3001::1 additional-paths send-receive`	Configure R3 to receive additional paths from the iBGP neighbor R2
`(config-router-af)#neighbor 3001::1 advertise additional-paths all`	Configure R2 to advertise all available paths to the iBGP neighbor R3
`(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 commit
`(config)#exit`	Exit the configure mode

© 2026 IP Infusion Inc. Proprietary

Validation

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

```
#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

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 commit
`(config)#exit`	Exit the configure mode

© 2026 IP Infusion Inc. Proprietary

Selection of all Additional Paths at the Neighbor Level

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)#neighbor 3001::2 additional-paths send`	Configure R2 to send additional paths to and receive additional paths from all iBGP neighbors
`(config-router-af)#neighbor 3001::2 advertise additional-paths 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 commit
`(config)#exit`	Exit the configure mode

Validation

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

```
#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
```

© 2026 IP Infusion Inc. Proprietary

Selection of Best 2 Additional Paths at AF Level

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`	Configure R2 to send additional paths to the iBGP neighbor R3
`(config-router-af)#bgp additional-paths select best 2`	Configure R2 to select best 2 out of all available paths 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 commit
`(config)#exit`	Exit the configure mode

© 2026 IP Infusion Inc. Proprietary

Selection of Best 2 Additional Paths at the Neighbor Level

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)#neighbor 3001::2 additional-paths send`	Configure R2 to send additional paths to the iBGP neighbor R3
`(config-router-af)#neighbor 3001::2 advertise additional-paths best 2`	Configure R2 to advertise best 2 out of all available paths to R3
`(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 commit
`(config)#exit`	Exit the configure mode

Validation

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
```

© 2026 IP Infusion Inc. Proprietary

```
      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

```
#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

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`	Configure R2 to send additional paths to the iBGP neighbor R3
`(config-router-af)#bgp additional-paths select best 3`	Configure R2 to select best 3 out of all available paths 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 commit
`(config)#exit`	Exit the configure mode

© 2026 IP Infusion Inc. Proprietary

Selection of Best 3 Additional Paths at the Neighbor Level

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)#neighbor 3001::2 additional-paths send`	Configure R2 to send additional paths to the iBGP neighbor R3
`(config-router-af)#neighbor 3001::2 advertise additional-paths best 3`	Configure R2 to advertise best 3 out of all available paths to R3
`(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 commit
`(config)#exit`	Exit the configure mode

Validation

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

```
#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
```

BGP Flowspec Support for IPv4

This section contains configurations of BGP Flowspec support for IPv4.

The BGP flow specification (flowspec) feature allows users to rapidly deploy and propagate filtering and policing functionality among a large number of BGP peer routers to mitigate the effects of a distributed denial-of-service (DDoS) attack over your network.

The BGP flowspec, specifically stated in here, allows users effectively to construct instructions to match a particular flow with source, destination, L4 parameters and packet specifics such as length, fragment and so on. Flowspec allows for a dynamic installation of an action at the border routers to either:

- Drop the traffic
- Inject it in a different VRF for analysis or
- Allow it, but police it at a specific defined rate

BGP Flowspec IPv4 Configuration

Topology

Figure 88. BGP Flowspec support for IPv4

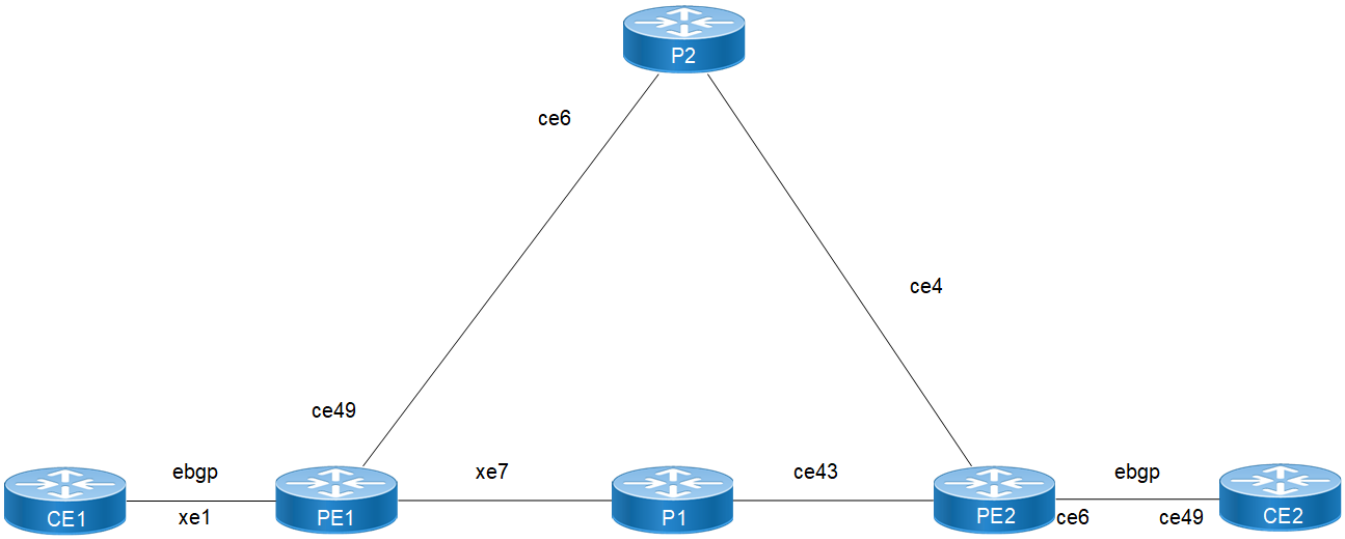

CE1 Configuration

`CE1#configure terminal`	Enter configure mode.
`CE1(config)# class-map type traffic match-all ocnos`	Configure class map with class name
`CE1(config-cmap-tr)#match destination-address ipv4 3.3.3.3/32`	Configure match with destination ip address
`CE1(config-cmap-tr)#match source-address ipv4 3.3.3.3/32`	Enable match with source ip address
`CE1(config-cmap-tr)# match destination-port 44-144`	Enable match with destination port

`CE1(config-cmap-tr)#exit`	Exit from class map traffic
`CE1(config)#commit`	Commits the configuration
`CE1(config)#class-map type traffic match-all ocnos2`	Configure class map with class name
`CE1(config-cmap-tr)# match destination-address ipv4 10.1.1.0/24`	Configure match with destination ip address
`CE1(config-cmap-tr)# match source-address ipv4 10.1.1.0/24`	Enable match with source ip address
`CE1(config-cmap-tr)# match source-port 22-122`	Enter match with source port
`CE1(config-cmap-tr)#exit`	Exit from class map traffic
`CE1(config)#commit`	Commits the configuration
`CE1(config)#policy-map type pbr CE1`	Configure policy map type pbr with policy name
`CE1(config-pmap-pbr)# class type traffic ocnos2`	Enter class map traffic class name
`CE1(config-pmap-pbr-c)# police rate 21000 bps`	Enable policy rate
`CE1(config-pmap-pbr-c)#exit`	Exit from polocy map traffic
`CE1(config-pmap-pbr)#exit`	Exit from policy map
`CE1(config)#policy-map type pbr PE1`	Configure policy map type pbr with policy name
`CE1(config-pmap-pbr)# class type traffic ocnos`	Enter class map traffic class name
`CE1(config-pmap-pbr-c)# police rate 4444444 bps`	Enable policy rate
`CE1(config-cmap-tr)#exit`	Exit from class map
`CE1(config)#commit`	Commits the configuration
`CE1(config)#flowspec`	Configure flowspec
`CE1(config-flowspec)#address-family ipv4`	Enter address-family mode with ipv4
`CE1(config-flowspec-af)#service-policy type pbr PE1`	Enable service plocy type pbr with policy name
`CE1(config-flowspec-af)#service-policy type pbr CE1`	Enable service plocy type pbr with policy name
`CE1(config-flowspec-af)#exit`	Exit from address family mode
`CE1(config-flowspec)#commit`	Commits the configuration
`CE1(config)#interface xe1`	Enter interface mode
`CE1(config-if)#ip address 10.1.1.1/24`	Configure ip address
`CE1(config-if)#exit`	Exit from interface mode
`CE1(config)#interface xe1.10`	Enter the sub interface mode
`CE1(config-if)# encapsulation dot1q 10`	Enable encapsulate dot 1q
`CE1(config-if)# ip address 66.1.1.1/24`	Configure ip address
`CE1(config-if)#exit`	Exit from interface mode
`CE1(config)#commit`	Commits the configuration
`CE1(config)#router bgp 100`	Configure bgp intstance
`CE1(config-router)#neighbor 10.1.1.2 remote-as 65000`	Enable neighbour ip address with remote-as 65000
`CE1(config-router)# neighbor 66.1.1.2 remote-as 65000`	Enable neighbour ip address with remote-as 65000

© 2026 IP Infusion Inc. Proprietary

`CE1(config-router)#address-family ipv4 unicast`	Enter address family ipv4 unicast mode
`CE1(config-router-af)#redistribute connected`	Enable redistribute connected
`CE1(config-router-af)#neighbor 10.1.1.2 activate`	Activate the neighbour
`CE1(config-router-af)#neighbor 66.1.1.2 activate`	Activate the neighbour
`CE1(config-router-af)#exit-address-family`	Exit from address family mode
`CE1(config-router)#address-family ipv4 flowspec`	Enter address family ipv4 flowspec mode
`CE1(config-router-af)#neighbor 10.1.1.2 activate`	Activate the neighbour
`CE1(config-router-af)#neighbor 66.1.1.2 activate`	Activate the neighbour
`CE1(config-router-af)#exit-address-family`	Exit from address family mode
`CE1(config-router)#exit`	Exit from router mode
`CE1(config)#commit`	Commits the configuration

PE1 Configuration

`PE1#configure terminal`	Enter configure mode.
`PE1(config)#ip vrf vrf100`	Enter vrf vrf name
`PE1(config-vrf)# rd 100:1`	Enter rd value
`PE1(config-vrf)# route-target both 100:1`	Enter router targeted as both
`PE1(config-vrf)#exit`	Exit from vrf mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#ip vrf vrf101`	Enter vrf vrf name
`PE1(config-vrf)# rd 101:1`	Enter rd value
`PE1(config-vrf)# route-target both 101:1`	Enter router targeted as both
`PE1(config-vrf)#exit`	Exit from vrf mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#qos enable`	Enable qos
`PE1(config)#qos statistics`	Enable qos statistics
`PE1(config)#class-map type traffic match-all clas1`	Configure class map with class name
`PE1(config-cmap-tr)# match destination-address ipv4 4.4.4.4/32`	Configure match with destination ip address
`PE1(config-cmap-tr)# match destination-port 400-499`	Enable match with destination-port
`PE1(config-cmap-tr)#exit`	Exit from class map traffic mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#policy-map type pbr pol1`	Configure policy map type pbr with policy name
`PE1(config-pmap-pbr)# class type traffic clas1`	Enable class type traffic with class map name
`PE1(config-pmap-pbr-c)# redirect ipv4 extcommunity rt 100:10`	Enable redirect ipv4 extcommunity with rt
`PE1(config-pmap-pbr-c)#exit`	Exit from policy map traffic

© 2026 IP Infusion Inc. Proprietary

`PE1(config-pmap-pbr)#exit`	Exit from policy map
`PE1(config)#commit`	Commits the configuration
`PE1(config)#flowspec`	Configure flowspec
`PE1(config-flowspec)# address-family ipv4`	Enter address-family mode with ipv4
`PE1(config-flowspec-af)# service-policy type pbr pol1`	Enable service policy type pbr with policy name
`PE1(config-flowspec-af)#exit`	Exit from address family mode
`PE1(config-flowspec)#exit`	Exit from flowspec mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#router ldp`	Configure router ldp
`PE1(config-router)# targeted-peer ipv4 3.3.3.3`	Configure targeted peer ipv4 address
`PE1(config-router-targeted-peer)# exit-targeted-peer-mode`	Exit from targeted peer mode
`PE1(config-router)#exit`	Exit from router mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#interface lo`	Enter interface mode
`PE1(config-if)# ip address 1.1.1.1/32 secondary`	Assign ip address as secondary
`PE1(config-if)# ip router isis 1`	Enable router isis instance
`PE1(config-if)#exit`	Exit from interface mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#interface ce51`	Enter interface mode
`PE1(config-if)# ip address 40.1.1.1/24`	Assign ip address
`PE1(config-if)# label-switching`	Enable label switching
`PE1(config-if)# ip router isis 1`	Enable router isis instance
`PE1(config-if)# enable-ldp ipv4`	Enable ldp ipv4
`PE1(config-if)#exit`	Exit from interface mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#interface xe1`	Enter interface mode
`PE1(config-if)# ip vrf forwarding vrf100`	Enable vrf forwarding with vrf name
`PE1(config-if)# ip address 10.1.1.2/24`	Assigh the ip address
`PE1(config-if)#exit`	Exit from interface mode
`PE1(config)#commit`	Commits the configuration
`PE1-7026(config)#interface xe1.10`	Enter the sub interface mode
`PE1-7026(config-if)# encapsulation dot1q 10`	Enable encapsulate dot 1q
`PE1-7026(config-if)# ip vrf forwarding vrf100`	Enable vrf forwarding under interface
`PE1-7026(config-if)# ip address 66.1.1.2/24`	Assigh the ip address
`PE1-7026(config-if)#exit`	Exit from interface mode

© 2026 IP Infusion Inc. Proprietary

`PE1(config)#interface xe7`	Enter the interface mode
`PE1(config-if)# ip address 20.1.1.1/24`	Enter ip address
`PE1(config-if)# label-switching`	Enable label switching
`PE1(config-if)# ip router isis 1`	Enable router isis instance
`PE1(config-if)# enable-ldp ipv4`	Enable ldp ipv4
`PE1(config-if)#exit`	Exit from interface mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#router isis 1`	Configure isis instance
`PE1(config-router)# is-type level-1-2`	Enable is-type level-1-2
`PE1(config-router)# metric-style wide`	Enable metric style wide
`PE1(config-router)# mpls traffic-eng router-id 1.1.1.1`	Enter router id
`PE1(config-router)# mpls traffic-eng level-1`	Enable traffic engg level-1
`PE1(config-router)# mpls traffic-eng level-2`	Enable traffic engg level-2
`PE1(config-router)# dynamic-hostname`	Enable hostname
`PE1(config-router)# bfd all-interfaces`	Enable bfd all the interfaces
`PE1(config-router)# net 49.0000.0000.0001.00`	Enter network id
`PE1(config-router)#exit`	Exit from isis instance
`PE1(config)#commit`	Commits the configuration
`PE1(config)#router bgp 65000`	Configure bgp instance
`PE1(config-router)# neighbor 3.3.3.3 remote-as 65000`	Enable neighbour address with remote-as 65000
`PE1(config-router)#neighbor 10.1.1.1 remote-as 100`	Enable neighbour address with remote-as 100
`PE1(config-router)#neighbor 3.3.3.3 update-source lo`	Enable neighbour address with update loppbacl
`PE1(config-router)# address-family ipv4 unicast`	Enter address family ipv4 unicast mode
`PE1(config-router-af)# neighbor 3.3.3.3 activate`	Activate the neighbour
`PE1(config-router-af)# neighbor 10.1.1.1 activate`	Activate the neighbour
`PE1(config-router-af)# exit-address-family`	Exit from address family mode
`PE1(config-router)# address-family vpnv4 unicast`	Enter address family vpnv4 mode
`PE1(config-router-af)# neighbor 3.3.3.3 activate`	Activate the neighbour
`PE1(config-router-af)# exit-address-family`	Exit from address family
`PE1(config-router)# address-family ipv4 flowspec`	Enter address family ipv4 flowspec mode
`PE1(config-router-af)# neighbor 3.3.3.3 activate`	Activate the neighbour
`PE1(config-router-af)# neighbor 10.1.1.1 activate`	Activate the neighbour
`PE1(config-router-af)# exit-address-family`	Exit from address family
`PE1(config-router)# address-family vpnv4 flowspec`	Enter address family vpnv4 flowpsec mode
`PE1(config-router-af)# neighbor 3.3.3.3 activate`	Activate the neighbour

© 2026 IP Infusion Inc. Proprietary

`PE1(config-router-af)# exit-address-family`	Exit from address family
`PE1(config-router)# address-family ipv4 vrf vrf100`	Enter address family ipv4 vrf mode
`PE1(config-router-af)# neighbor 66.1.1.1 remote-as 100`	Enter neighbor ip add with remote-as 100
`PE1(config-router-af)# neighbor 66.1.1.1 activate`	Activate the neighbour
`PE1(config-router)# address-family ipv4 flowspec vrf vrf100`	Enter address family ipv4 vrf flowspec mode
`PE1(config-router-af)# neighbor 66.1.1.1 activate`	Activate the neighbour
`PE1(config-router-af)# exit-address-family`	Exit from address family
`PE1(config-router)#exit`	Exit from router mode
`PE1(config)#commit`	Commits the configuration

P1 Configuration

`P1#configure terminal`	Enter configure mode.
`P1(config)#router ldp`	Enter router ldp
`P1(config-router)#exit`	Exit from ldp mode
`P1(config)#commit`	Commits the configuration
`P1(config)#interface lo`	Enter interface mode
`P1(config-if)#ip address 2.2.2.2/32 secondary`	Assigned ip add as secondary
`P1(config-if)# ip router isis 1`	Enter isis instance
`P1(config-if)#exit`	Exit from interface mode
`P1(config)#commit`	Commits the configuration
`P1(config)#interface ce43`	Enter interface mode
`P1(config-if)# speed 40g`	Set the speed
`P1(config-if)# ip address 11.1.1.2/24`	Assign the ip add
`P1(config-if)# label-switching`	Enable label switching
`P1(config-if)# ip router isis 1`	Enter router isis instance
`P1(config-if)# enable-ldp ipv4`	Enable ldp ipv4
`P1(config-if)#exit`	Exit from interface mode
`P1(config)#commit`	Commits the configuration
`P1(config)#interface xe7`	Enter interface mode
`P1(config-if)# ip address 20.1.1.2/24`	Assign the ip add
`P1(config-if)# label-switching`	Enable label switching
`P1(config-if)# ip router isis 1`	Enter router isis instance
`P1(config-if)# enable-ldp ipv4`	Enable ldp ipv4
`P1(config-if)#exit`	Exit from interface mode
`P1(config)#commit`	Commits the configuration

© 2026 IP Infusion Inc. Proprietary

P1(config)#router isis 1	Configure isis instance
P1(config-router)# is-type level-1-2	Configure is-type-level-1-2
P1(config-router)# metric-style wide	Configure metric style wide
P1(config-router)# mpls traffic-eng router-id 2.2.2.2	Configure router id
P1(config-router)# mpls traffic-eng level-1	Configure traffic engg level-1
P1(config-router)# mpls traffic-eng level-2	Configure traffic engg level-2
P1(config-router)# dynamic-hostname	Configure hostname
P1(config-router)# bfd all-interfaces	Enable bfd all interfaces
P1(config-router)# net 49.0000.0000.0002.00	Configure network id
P1(config-router)#exit	Exit from isis instance mode
P1(config)#commit	Commits the configuration

P2 Configuration

P2#configure terminal	Enter configure mode.
P2(config)#router ldp	Enter router ldp
P2(config-router)#exit	Exit from ldp mode
P2(config)#commit	Commits the configuration
P2(config)#interface lo	Enter the interface mode
P2(config-if)# ip address 6.6.6.6/32 secondary	Assign the ip address secondary
P2(config-if)# ip router isis 1	Enter router isis instance
P2(config-if)#exit	Exit from interface mode
P2(config)#commit	Commits the configuration
P2(config)#interface ce53	Enter the interface mode
P2(config-if)# ip address 50.1.1.2/24	Assign the ip add
P2(config-if)# label-switching	Enable label switching
P2(config-if)# ip router isis 1	Enter router isis instance
P2(config-if)# enable-ldp ipv4	Enable ldp ipv4
P2(config-if)#exit	Exit from interface mode
P2(config)#commit	Commits the configuration
P2(config)#interface ce51	Enter the interface mode
P2(config-if)# ip address 40.1.1.2/24	Assign the ip add
P2(config-if)# label-switching	Enable label switching
P2(config-if)# ip router isis 1	Enter router isis instance
P2(config-if)# enable-ldp ipv4	Enable ldp ipv4
P2(config-if)#exit	Exit from interface mode
P2(config)#commit	Commits the configuration

© 2026 IP Infusion Inc. Proprietary

`P2(config)#router isis 1`	Configure isis instance
`P2(config-router)# is-type level-1-2`	Configure is-type-level-1-2
`P2(config-router)# metric-style wide`	Configure metric style wide
`P2(config-router)# mpls traffic-eng router-id 6.6.6.6`	Configure router id
`P2(config-router)# mpls traffic-eng level-1`	Configure traffic engg level-1
`P2(config-router)# mpls traffic-eng level-2`	Configure traffic engg level-2
`P2(config-router)# dynamic-hostname`	Configure hostname
`P2(config-router)# bfd all-interfaces`	Enable bfd all interfaces
`P2(config-router)# net 49.0000.0000.0006.00`	Configure network id
`P2(config-router)#exit`	Exit from isis instance mode
`P2(config)#commit`	Commits the configuration

PE2 Configuration

`PE2#configure terminal`	Enter configure mode.
`PE2(config)#router ldp`	Configure router ldp
`PE2(config-router)# targeted-peer ipv4 1.1.1.1`	Configure targeted peer ipv4 address
`PE2(config-router-targeted-peer)#exit-targeted-peer-mode`	Exit from targeted peer mode
`PE2(config-router)#exit`	Exit from router mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#ip vrf vrf100`	Enter vrf vrf name
`PE2(config-vrf)# rd 100:1`	Enter rd value
`PE2(config-vrf)# route-target both 100:1`	Enter router targeted as both
`PE2(config-vrf)#exit`	Exit from vrf mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#ip vrf vrf101`	Enter vrf vrf name
`PE2(config-vrf)# rd 101:1`	Enter rd value
`PE2(config-vrf)# route-target both 101:1`	Enter router targeted as both
`PE2(config-vrf)#exit`	Exit from vrf mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#class-map type traffic match-all cl1`	Configure class map with class name
`PE2(config-cmap-tr)# match destination-address ipv4 1.1.1.1/32`	Configure match with destination ip address
`PE2(config-cmap-tr)# match source-address ipv4 3.3.3.3/32`	Enable match with source ip address
`PE2(config-cmap-tr)# match destination-port 300-399`	Enable match with destination-port
`PE2(config-cmap-tr)#exit`	Exit from class map traffic mode

© 2026 IP Infusion Inc. Proprietary

`PE2(config)#commit`	Commits the configuration
`PE2(config)#policy-map type pbr pol1`	Configure policy map type pbr with policy name
`PE2(config-pmap-pbr)# class type traffic cl1`	Enable class type traffic with class map name
`PE2(config-pmap-pbr-c)# police rate 1111111 bps`	Enable policy rate
`PE2(config-pmap-pbr-c)#exit`	Exit from polocy map traffic
`PE2(config-pmap-pbr)#exit`	Exit from policy map
`PE2(config)#commit`	Commits the configuration
`PE2(config)#flowspec`	Configure flowspec
`PE2(config-flowspec)# address-family ipv4`	Enter address-family mode with ipv4
`PE2(config-flowspec-af)# service-policy type pbr pol1`	Enable service plocy type pbr with policy name
`PE2(config-flowspec-af)#exit`	Exit from address family mode
`PE2(config-flowspec)#exit`	Exit from flowspec mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#interface lo`	Enter the interface mode
`PE2(config-if)# ip address 3.3.3.3/32 secondary`	Assigne the ip add secondary
`PE2(config-if)# ip router isis 1`	Enter router isis instance
`PE2(config-if)#exit`	Exit from interface mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#interface ce4`	Enter the interface mode
`PE2(config-if)# ip address 50.1.1.1/24`	Assign the ip add
`PE2(config-if)# label-switching`	Enable label switching
`PE2(config-if)# ip router isis 1`	Enter router isis instance
`PE2(config-if)# enable-ldp ipv4`	Enable ldp ipv4
`PE2(config-if)#exit`	Exit from interface mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#interface ce43`	Enter the interface mode
`PE2(config-if)# speed 40g`	Set the interface speed
`PE2(config-if)# ip address 11.1.1.1/24`	Assign the ip add
`PE2(config-if)# label-switching`	Enable label switching
`PE2(config-if)# ip router isis 1`	Enter router isis instance
`PE2(config-if)#enable-ldp ipv4`	Enable ldp ipv4
`PE2(config-if)#exit`	Exit from interface mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#interface ce6`	Enter the interface mode
`PE2(config-if)# ip address 30.1.1.2/24`	Assign the ip add
`PE2(config-if)#exit`	Exit from interface mode

© 2026 IP Infusion Inc. Proprietary

`PE2(config)#commit`	Commits the configuration
`PE2-7023(config)#interface ce6.10`	Enter the sub interface mode
`PE2-7023(config-if)# encapsulation dot1q 10`	Enable encapsulate dot 1q
`PE2-7023(config-if)# ip vrf forwarding vrf100`	Enable vrf forwarding vrf under interface
`PE2-7023(config-if)# ip address 55.1.1.1/24`	Assign the ip add
`PE2-7023(config-if)#exit`	Exit from interface mode
`PE2(config)#router isis 1`	Configure isis instance
`PE2(config-router)# is-type level-1-2`	Configure is-type-level-1-2
`PE2(config-router)# metric-style wide`	Configure metric style wide
`PE2(config-router)# mpls traffic-eng router-id 3.3.3.3`	Configure router id
`PE2(config-router)# mpls traffic-eng level-1`	Configure traffic engg level-1
`PE2(config-router)# mpls traffic-eng level-2`	Configure traffic engg level-2
`PE2(config-router)# dynamic-hostname`	Configure hostname
`PE2(config-router)# bfd all-interfaces`	Enable bfd all interfaces
`PE2(config-router)# net 49.0000.0000.0003.00`	Configure network id
`PE2(config-router)#exit`	Exit from isis instance mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#router bgp 65000`	Configure bgp instance
`PE2(config-router)# neighbor 1.1.1.1 remote-as 65000`	Enable neighbour address with remote-as 65000
`PE2(config-router)# neighbor 30.1.1.1 remote-as 200`	Enable neighbour address with remote-as 200
`PE2(config-router)# neighbor 1.1.1.1 update-source lo`	Enable neighbour address with update loppbacl
`PE2(config-router)# address-family ipv4 unicast`	Enter address family ipv4 unicast mode
`PE2(config-router-af)# neighbor 1.1.1.1 activate`	Activate the neighbour
`PE2(config-router-af)# neighbor 30.1.1.1 activate`	Activate the neighbour
`PE2(config-router-af)# exit-address-family`	Exit from address family mode
`PE2(config-router)# address-family vpnv4 unicast`	Enter address family vpnv4 mode
`PE2(config-router-af)# neighbor 1.1.1.1 activate`	Activate the neighbour
`PE2(config-router-af)# exit-address-family`	Exit from address family
`PE2(config-router)# address-family ipv4 flowspec`	Enter address family ipv4 flowspec mode
`PE2(config-router-af)# neighbor 1.1.1.1 activate`	Activate the neighbour
`PE2(config-router-af)# neighbor 30.1.1.1 activate`	Activate the neighbour
`PE2(config-router-af)# exit-address-family`	Exit from address family
`PE2(config-router)# address-family vpnv4 flowspec`	Enter address family vpnv4 flowpsec mode
`PE2(config-router-af)# neighbor 1.1.1.1 activate`	Activate the neighbour
`PE2(config-router-af)# exit-address-family`	Exit from address family

© 2026 IP Infusion Inc. Proprietary

`PE2(config-router)# address-family ipv4 vrf vrf100`	Enter address family ipv4 vrf mode
`PE2(config-router-af)# neighbor 55.1.1.2 remote-as 200`	Enter neighbor ip add with remote-as 100
`PE2(config-router-af)# neighbor 55.1.1.2 activate`	Activate the neighbour
`PE2(config-router-af)# exit-address-family`	Exit from address family
`PE2(config-router)# address-family ipv4 flowspec vrf vrf100`	
`PE2(config-router-af)# neighbor 55.1.1.2 activate`	Activate the neighbour
`PE2(config-router-af)# exit-address-family`	Exit from address family
`PE2(config-router)#exit`	Exit from router mode
`PE2(config)#commit`	Commits the configuration

CE2 Configuration

`CE2#configure terminal`	Enter configure mode.
`CE2(config)#class-map type traffic match-all clas1`	Configure class map with class name
`CE2(config-cmap-tr)# match destination-address ipv4 4.4.4.4/32`	Configure match with destination ip address
`CE2(config-cmap-tr)# match source-address ipv4 5.5.5.5/32`	Enable match with source ip address
`CE2(config-cmap-tr)# match destination-port 1300-1400`	Enable match with destination-port
`CE2(config-cmap-tr)#exit`	Exit from class map traffic mode
`CE2(config)#commit`	Commits the configuration
`CE2(config)#class-map type traffic match-all clas2`	Configure class map with class name
`CE2(config-cmap-tr)# match destination-address ipv4 30.1.1.0/24`	Configure match with destination ip address
`CE2(config-cmap-tr)# match source-address ipv4 30.1.1.0/24`	Enable match with source ip address
`CE2(config-cmap-tr)# match source-port 500-600`	Enable match with source-port
`CE2(config-cmap-tr)# match tcp-flag 21 bit-mask 3f`	Enable match with tcp flag bit
`CE2(config-cmap-tr)# match dscp 21`	Enable match with dscp
`CE2(config-cmap-tr)#exit`	Exit from class map traffic mode
`CE2(config)#commit`	Commits the configuration
`CE2(config)#policy-map type pbr pol1`	Configure policy map type pbr with policy name
`CE2(config-pmap-pbr)#class type traffic clas1`	Enable class type traffic with class map name
`CE2(config-pmap-pbr-c)#police rate 100000 bps`	Enable policy rate
`CE2(config-pmap-pbr-c)#exit`	Exit from polocy map traffic
`CE2(config-pmap-pbr)#exit`	Exit from policy map
`CE2(config)#commit`	Commits the configuration
`CE2(config)#policy-map type pbr pol2`	Configure policy map type pbr with policy name

`CE2(config-pmap-pbr)# class type traffic clas2`	Enable class type traffic with class map name
`CE2(config-pmap-pbr-c)# police rate 500000 bps`	Enable policy rate
`CE2(config-pmap-pbr-c)# set dscp 1e`	Exit from polocy map traffic
`CE2(config-pmap-pbr-c)#exit`	Exit from policy map
`CE2(config-pmap-pbr)#exit`	Commits the configuration
`CE2(config)#commit`	Configure policy map type pbr with policy name
`CE2(config)#flowspec`	Configure flowspec
`CE2(config-flowspec)# address-family ipv4`	Enter address-family mode with ipv4
`CE2(config-flowspec-af)# service-policy type pbr pol1`	Enable service plocy type pbr with policy name
`CE2(config-flowspec-af)# service-policy type pbr pol2`	Enable service plocy type pbr with policy name
`CE2(config-flowspec-af)#exit`	Exit from address family mode
`CE2(config-flowspec)#exit`	Exit from flowspec mode
`CE2(config)#commit`	Commits the configuration
`CE2(config)#interface lo`	Enter interface mode
`CE2(config-if)# ip address 5.5.5.5/32 secondary`	Assign the ip address secondary
`CE2(config-if)# ip router isis 1`	Enter router isis instance
`CE2(config-if)#exit`	Exit from interface mode
`CE2(config)#commit`	Commits the configuration
`CE2(config)#interface ce49`	Enter interface mode
`CE2(config-if)# ip address 30.1.1.1/24`	Enter interface mode
`CE2(config-if)#exit`	Exit from interface mode
`CE2(config)#commit`	Commits the configuration
`CE2(config)#router bgp 200`	Configure bgp instance
`CE2(config-router)# neighbor 30.1.1.2 remote-as 65000`	Enable neighbour address with remote-as 65000
`CE2(config-router)# neighbor 55.1.1.1 remote-as 65000`	Enable neighbour address with remote-as 200
`CE2(config-router)# address-family ipv4 unicast`	Enter address family ipv4 unicast mode
`CE2(config-router-af)# redistribute connected`	Enter redistribute connected
`CE2(config-router-af)# neighbor 30.1.1.2 activate`	Activate the neighbour
`CE2(config-router-af)# neighbor 55.1.1.1 activate`	Activate the neighbour
`CE2(config-router-af)# exit-address-family`	Exit from address family mode
`CE2(config-router)# address-family ipv4 flowspec`	Enter address family ipv4 flowspec mode
`CE2(config-router-af)# neighbor 30.1.1.2 activate`	Activate the neighbour
`CE2(config-router-af)# neighbor 55.1.1.1 activate`	Activate the neighbour
`CE2(config-router-af)# exit-address-family`	Exit from address family

© 2026 IP Infusion Inc. Proprietary

`CE2(config-router)#exit`	Exit from router bgp mode
`CE2(config)#commit`	Commits the configuration

BGP Flowspec IPv4 Validation

CE1

The CE1 node configurations are verified using the show commands listed below.

```
CE1#sh ip bgp
BGP table version is 5, local router ID is 4.4.4.4
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
*>   4.4.4.4/32       0.0.0.0              0        100       32768  ?
*>   5.5.5.5/32       66.1.1.2             0        100       0   65000 200 ?
*>   10.1.1.0/24      0.0.0.0              0        100       32768  ?
*>   30.1.1.0/24      66.1.1.2             0        100       0   65000 200 ?
*>   55.1.1.0/24      66.1.1.2             0        100       0   65000 ?
*>   66.1.1.0/24      0.0.0.0              0        100       32768  ?
*                     66.1.1.2             0        100       0   65000 ?
*>   100.1.1.0/24     0.0.0.0              0        100       32768  ?
*>   200.1.1.0        66.1.1.2             0        100       0   65000 200 ?

Total number of prefixes 8
CE1#
```

```
CE1#show ip bgp neighbors
BGP neighbor is 10.1.1.2, remote AS 65000, local AS 100, external link
  BGP version 4, local router ID 4.4.4.4, remote router ID 1.1.1.1
  BGP state = Established, up for 02:20:12
  Last read 00:00:07, 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 IPv4 Flow Specification: advertised and received
  Received 345 messages, 0 notifications, 0 in queue
  Sent 358 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 5, neighbor version 5
  Index 0, Offset 0, Mask 0x1
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  8 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 8, neighbor version 8
  Index 0, Offset 0, Mask 0x1
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes

 Connections established 1; dropped 0
Local host: 10.1.1.1, Local port: 179
Foreign host: 10.1.1.2, Foreign port: 59044
Nexthop: 10.1.1.1
Nexthop global: ::
```

```
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 66.1.1.2, remote AS 65000, local AS 100, external link
  BGP version 4, local router ID 4.4.4.4, remote router ID 66.1.1.2
  BGP state = Established, up for 02:20:12
  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
    Address family IPv4 Flow Specification: advertised and received
  Received 348 messages, 0 notifications, 0 in queue
  Sent 343 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 5, neighbor version 5
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  5 accepted prefixes
  4 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 8, neighbor version 8
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 66.1.1.1, Local port: 179
Foreign host: 66.1.1.2, Foreign port: 54954
Nexthop: 66.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
```

```
CE1#show ip bgp flowspec summary
BGP router identifier 4.4.4.4, local AS number 100
BGP table version is 8
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
10.1.1.2                 4 65000  346        360       8      0      0  02:20:46               4
66.1.1.2                 4 65000  349        344       8      0      0  02:20:46               2

Total number of neighbors 2

Total number of Established sessions 2
CE1#

CE1#show ip bgp ipv4 flowspec
BGP table version is 8, local router ID is 4.4.4.4
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
*>   flowspec/136     0.0.0.0               0        100       32768  i
*>   flowspec/200     66.1.1.2              0        100       0   65000 200 i
*                     10.1.1.2              0        100       0   65000 200 i
*    flowspec/152     10.1.1.2              0        100       0   65000 i
```

```
*>   flowspec/152     0.0.0.0              0        100       32768  i
*    flowspec/48      10.1.1.2             0        100       0   65000 i
*>   flowspec/152     66.1.1.2             0        100       0   65000 200 i
*                     10.1.1.2             0        100       0   65000 200 i

Total number of prefixes 6
CE1#

CE1#show ip bgp ipv4 flowspec vrf all
BGP table version is 8, local router ID is 4.4.4.4
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
*>   flowspec/136     0.0.0.0              0        100       32768  i
*>   flowspec/200     66.1.1.2             0        100       0   65000 200 i
*                     10.1.1.2             0        100       0   65000 200 i
*    flowspec/152     10.1.1.2             0        100       0   65000 i
*>   flowspec/152     0.0.0.0              0        100       32768  i
*    flowspec/48      10.1.1.2             0        100       0   65000 i
*>   flowspec/152     66.1.1.2             0        100       0   65000 200 i
*                     10.1.1.2             0        100       0   65000 200 i

Total number of prefixes 6
CE1#
```

```
CE1#show ip bgp ipv4 flowspec vrf all detail
BGP table version is 8, local router ID is 4.4.4.4
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
*>   flowspec/136     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>   flowspec/200     66.1.1.2             0        100       0   65000 200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*                     10.1.1.2             0        100       0   65000 200 i
```

```
  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*    flowspec/152     10.1.1.2             0        100       0   65000 i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*>   flowspec/152     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 3.3.3.3/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*    flowspec/48      10.1.1.2             0        100       0   65000 i

  FlowSpec:
    destination ip address: 4.4.4.4/32

  FlowSpec Actions:
    rt_redirect: as_type: as 100:10 vrf_id:0 fib_id:0

*>   flowspec/152     66.1.1.2             0        100       0   65000 200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400

  FlowSpec Actions:
    traffic_rate: 100000 bps

*                    10.1.1.2             0        100       0   65000 200 i
  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 6
CE1#
```

```
CE1#show ip bgp ipv4 flowspec detail
BGP table version is 8, local router ID is 4.4.4.4
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
*>   flowspec/136     0.0.0.0               0        100       32768  i

  FlowSpec:
    destination ip address: 10.1.1.0/24
```

```
    source ip address: 10.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>   flowspec/200     66.1.1.2             0        100       0   65000 200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*                    10.1.1.2             0        100       0   65000 200 i
  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*    flowspec/152     10.1.1.2             0        100       0   65000 i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*>   flowspec/152     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 3.3.3.3/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*    flowspec/48      10.1.1.2             0        100       0   65000 i

  FlowSpec:
    destination ip address: 4.4.4.4/32

  FlowSpec Actions:
    rt_redirect: as_type: as 100:10 vrf_id:0 fib_id:0

*>   flowspec/152     66.1.1.2             0        100       0   65000 200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
```

```
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400

  FlowSpec Actions:
    traffic_rate: 100000 bps

*                    10.1.1.2             0        100       0   65000 200 i
  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 6
CE1#
```

PE1

The PE1 node configurations are verified using the show commands listed below.

```
PE1#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
P1             xe7         3417.ebe4.af08      Up     27        L1   IS-IS
                                               Up     27        L2   IS-IS
P2             ce51        3c2c.99c7.07ad      Up     19        L1   IS-IS
                                               Up     19        L2   IS-IS
PE1#

PE1#show ip bgp flowspec summary
BGP router identifier 1.1.1.1, local AS number 65000
BGP table version is 17
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
3.3.3.3                 4 65000  378        368      17      0      0  02:27:00               3
10.1.1.1                4   100  373        361      17      0      0  02:27:00               2

Total number of neighbors 2

Total number of Established sessions 2
BGP router identifier 66.1.1.2, local AS number 65000
BGP VRF vrf100 Route Distinguisher: 100:1
BGP table version is 1
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
66.1.1.1                4   100  358        363       1      0      0  02:27:00               2

Total number of neighbors 1

Total number of Established sessions 1
PE1#

PE1#sh ip bgp ipv4 flowspec
BGP table version is 17, local router ID is 1.1.1.1
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
*>   flowspec/136     10.1.1.1              0        100       0   100 i
*>i  flowspec/200     3.3.3.3               0        100       0   200 i
*>i  flowspec/152     3.3.3.3               0        100       0    i
*    flowspec/152     10.1.1.1              0        100       0   100 i
*>   flowspec/48      0.0.0.0               0        100       32768  i
*>i  flowspec/152     3.3.3.3               0        100       0   200 i

Total number of prefixes 6
PE1#

PE1#show ip bgp vpnv4 all
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
Route Distinguisher: 100:1 (Default for VRF vrf100)
*> l 4.4.4.4/32       66.1.1.1              0        100       0   100 ?
*>i  5.5.5.5/32       3.3.3.3               0        100       0   200 ?
*> l 10.1.1.0/24      66.1.1.1              0        100       0   100 ?
*>i  30.1.1.0/24      3.3.3.3               0        100       0   200 ?
*>i  55.1.1.0/24      3.3.3.3               0        100       0    ?
*> l 66.1.1.0/24      0.0.0.0               0        100       32768  ?
*                     66.1.1.1              0        100       0   100 ?
*> l 100.1.1.0/24     66.1.1.1              0        100       0   100 ?
*>i  200.1.1.0        3.3.3.3               0        100       0   200 ?
 Announced routes count = 5
 Accepted routes count = 4
Route Distinguisher: 100:1
*>i  5.5.5.5/32       3.3.3.3               0        100       0   200 ?
*>i  30.1.1.0/24      3.3.3.3               0        100       0   200 ?
*>i  55.1.1.0/24      3.3.3.3               0        100       0    ?
*>i  200.1.1.0        3.3.3.3               0        100       0   200 ?
 Announced routes count = 0
 Accepted routes count = 4
PE1#

PE1#show ip bgp vpnv4 flowspec summary
BGP router identifier 1.1.1.1, local AS number 65000
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
3.3.3.3                  4 65000  384        375       5      0      0  02:30:02              2

Total number of neighbors 1

Total number of Established sessions 1
PE1#

PE1#show ip bgp vpnv4 flowspec vrf vrf100
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
Route Distinguisher: 100:1 (Default for VRF vrf100)
*>   flowspec/136     66.1.1.1              0        100       0   100 i
*>i  flowspec/200     3.3.3.3               0        100       0   200 i
*    flowspec/152     66.1.1.1              0        100       0   100 i
*>i  flowspec/152     3.3.3.3               0        100       0   200 i
 Announced routes count = 2
 Accepted routes count = 2
```

```
PE1#

PE1#show ip bgp vpnv4 flowspec vrf vrf100 detail
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
Route Distinguisher: 100:1 (Default for VRF vrf100)
*>   flowspec/136     66.1.1.1              0        100        0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>i  flowspec/200     3.3.3.3               0        100        0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*    flowspec/152     66.1.1.1              0        100        0   100 i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 3.3.3.3/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>i  flowspec/152     3.3.3.3               0        100        0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400

  FlowSpec Actions:
    traffic_rate: 100000 bps

 Announced routes count = 2
 Accepted routes count = 2
PE1#
```

```
PE1#sh ip bgp ipv4 flowspec detail
BGP table version is 17, local router ID is 1.1.1.1
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
*>   flowspec/136     10.1.1.1              0        100       0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>i  flowspec/200     3.3.3.3               0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*>i  flowspec/152     3.3.3.3               0        100       0    i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*    flowspec/152     10.1.1.1              0        100       0   100 i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 3.3.3.3/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>   flowspec/48      0.0.0.0               0        100       32768  i

  FlowSpec:
    destination ip address: 4.4.4.4/32

  FlowSpec Actions:
    rt_redirect: as_type: as 100:10 vrf_id:0 fib_id:0
```

```
*>i  flowspec/152     3.3.3.3              0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 6
PE1#

PE1#show ip bgp ipv4 flowspec vrf all
BGP table version is 1, local router ID is 66.1.1.2
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

BGP Route Table for VRF vrf100
*>   flowspec/136     66.1.1.1             0        100       0   100 i
*>i  flowspec/200     3.3.3.3              0        100       0   200 i
*    flowspec/152     66.1.1.1             0        100       0   100 i
*>i  flowspec/152     3.3.3.3              0        100       0   200 i

Total number of prefixes 4
*>   flowspec/136     10.1.1.1             0        100       0   100 i
*>i  flowspec/200     3.3.3.3              0        100       0   200 i
*>i  flowspec/152     3.3.3.3              0        100       0    i
*    flowspec/152     10.1.1.1             0        100       0   100 i
*>   flowspec/48      0.0.0.0              0        100       32768  i
*>i  flowspec/152     3.3.3.3              0        100       0   200 i

Total number of prefixes 6
PE1#
PE1#show ip bgp ipv4 flowspec vrf all detail
BGP table version is 1, local router ID is 66.1.1.2
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

BGP Route Table for VRF vrf100
*>   flowspec/136     66.1.1.1             0        100       0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>i  flowspec/200     3.3.3.3              0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    TCP/UDP source port:
```

```
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*    flowspec/152     66.1.1.1             0        100       0   100 i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 3.3.3.3/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>i  flowspec/152     3.3.3.3              0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 4
*>   flowspec/136     10.1.1.1             0        100       0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>i  flowspec/200     3.3.3.3              0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e
```

© 2026 IP Infusion Inc. Proprietary

```
*>i  flowspec/152     3.3.3.3              0        100       0    i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*    flowspec/152     10.1.1.1             0        100       0   100 i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 3.3.3.3/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>   flowspec/48      0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 4.4.4.4/32

  FlowSpec Actions:
    rt_redirect: as_type: as 100:10 vrf_id:0 fib_id:0

*>i  flowspec/152     3.3.3.3              0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 6
PE1#
```

```
PE1#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
PE1.00-00           * 0x00000013   0x9F2C        1079              0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     PE1
  IP Address:   20.1.1.1
  Router ID:    1.1.1.1
  Metric:   10         IS-Extended PE1.02
    IPv4 Interface Address: 20.1.1.1
    Neighbor IP Address: 20.1.1.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
```

```
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.03
    IPv4 Interface Address: 40.1.1.1
    Neighbor IP Address: 40.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 20.1.1.0/24
  Metric:   0          IP-Extended 1.1.1.1/32
  Metric:   10         IP-Extended 40.1.1.0/24
PE1.02-00           * 0x0000000E   0x704C        640               0/0/0
  Metric:   0          IS-Extended PE1.00
  Metric:   0          IS-Extended P1.00
PE1.03-00           * 0x00000003   0xCFF2        743               0/0/0
  Metric:   0          IS-Extended PE1.00
  Metric:   0          IS-Extended P2.00
P1.00-00              0x00000013   0x6CF5        1159              0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     P1
  IP Address:   2.2.2.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended PE2.02
    IPv4 Interface Address: 11.1.1.2
    Neighbor IP Address: 11.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.02
    IPv4 Interface Address: 20.1.1.2
    Neighbor IP Address: 20.1.1.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 20.1.1.0/24
  Metric:   10         IP-Extended 2.2.2.2/32
```

```
  Metric:   10         IP-Extended 11.1.1.0/24
PE2.00-00             0x00000013   0x2F15        1167              0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     PE2
  IP Address:   11.1.1.1
  Router ID:    3.3.3.3
  Metric:   10         IS-Extended PE2.02
    IPv4 Interface Address: 11.1.1.1
    Neighbor IP Address: 11.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE2.01
    IPv4 Interface Address: 50.1.1.1
    Neighbor IP Address: 50.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   0          IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 11.1.1.0/24
PE2.01-00             0x0000000D   0xCBEA        542               0/0/0
  Metric:   0          IS-Extended PE2.00
  Metric:   0          IS-Extended P2.00
PE2.02-00             0x0000000E   0x7246        503               0/0/0
  Metric:   0          IS-Extended PE2.00
  Metric:   0          IS-Extended P1.00
P2.00-00              0x00000011   0x788C        1147              0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     P2
  IP Address:   6.6.6.6
  Router ID:    6.6.6.6
  Metric:   10         IS-Extended PE2.01
    IPv4 Interface Address: 50.1.1.2
    Neighbor IP Address: 50.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.03
```

```
    IPv4 Interface Address: 40.1.1.2
    Neighbor IP Address: 40.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 6.6.6.6/32
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IP-Extended 50.1.1.0/24

IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
PE1.00-00           * 0x00000017   0x40F9        1079              0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     PE1
  IP Address:   20.1.1.1
  Router ID:    1.1.1.1
  Metric:   10         IS-Extended PE1.02
    IPv4 Interface Address: 20.1.1.1
    Neighbor IP Address: 20.1.1.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.03
    IPv4 Interface Address: 40.1.1.1
    Neighbor IP Address: 40.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 20.1.1.0/24
  Metric:   0          IP-Extended 1.1.1.1/32
  Metric:   20         IP-Extended 2.2.2.2/32
  Metric:   20         IP-Extended 11.1.1.0/24
  Metric:   20         IP-Extended 3.3.3.3/32
  Metric:   20         IP-Extended 50.1.1.0/24
  Metric:   20         IP-Extended 6.6.6.6/32
  Metric:   10         IP-Extended 40.1.1.0/24
PE1.02-00           * 0x0000000E   0x704C        640               0/0/0
  Metric:   0          IS-Extended PE1.00
  Metric:   0          IS-Extended P1.00
PE1.03-00           * 0x00000003   0xCFF2        743               0/0/0
```

```
  Metric:   0          IS-Extended PE1.00
  Metric:   0          IS-Extended P2.00
P1.00-00              0x00000019   0xCCF2        1159              0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     P1
  IP Address:   2.2.2.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended PE2.02
    IPv4 Interface Address: 11.1.1.2
    Neighbor IP Address: 11.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.02
    IPv4 Interface Address: 20.1.1.2
    Neighbor IP Address: 20.1.1.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 20.1.1.0/24
  Metric:   10         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 11.1.1.0/24
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   20         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 1.1.1.1/32
  Metric:   30         IP-Extended 6.6.6.6/32
  Metric:   20         IP-Extended 40.1.1.0/24
PE2.00-00             0x00000016   0x1BA1        1167              0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     PE2
  IP Address:   11.1.1.1
  Router ID:    3.3.3.3
  Metric:   10         IS-Extended PE2.02
    IPv4 Interface Address: 11.1.1.1
    Neighbor IP Address: 11.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE2.01
```

```
    IPv4 Interface Address: 50.1.1.1
    Neighbor IP Address: 50.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   0          IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 11.1.1.0/24
  Metric:   20         IP-Extended 2.2.2.2/32
  Metric:   20         IP-Extended 20.1.1.0/24
  Metric:   20         IP-Extended 1.1.1.1/32
  Metric:   20         IP-Extended 6.6.6.6/32
  Metric:   20         IP-Extended 40.1.1.0/24
PE2.01-00             0x0000000D   0xCBEA        542               0/0/0
  Metric:   0          IS-Extended PE2.00
  Metric:   0          IS-Extended P2.00
PE2.02-00             0x0000000E   0x7246        503               0/0/0
  Metric:   0          IS-Extended PE2.00
  Metric:   0          IS-Extended P1.00
P2.00-00              0x00000014   0x3E72        1147              0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     P2
  IP Address:   6.6.6.6
  Router ID:    6.6.6.6
  Metric:   10         IS-Extended PE2.01
    IPv4 Interface Address: 50.1.1.2
    Neighbor IP Address: 50.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.03
    IPv4 Interface Address: 40.1.1.2
    Neighbor IP Address: 40.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 6.6.6.6/32
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 1.1.1.1/32
```

```
  Metric:   30         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   20         IP-Extended 11.1.1.0/24
  Metric:   20         IP-Extended 20.1.1.0/24

PE1#
```

```
PE1#show ip bgp neighbors
BGP neighbor is 3.3.3.3, remote AS 65000, local AS 65000, internal link
  BGP version 4, local router ID 1.1.1.1, remote router ID 3.3.3.3
  BGP state = Established, up for 02:26:21
  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
    Address family VPNv4 Unicast: advertised and received
    Address family IPv4 Flow Specification: advertised and received
    Address family VPNv4 Flow Specification: advertised and received
  Received 376 messages, 0 notifications, 0 in queue
  Sent 367 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 6, neighbor version 6
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: VPNv4 Unicast  BGP table version 4, neighbor version 4
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 17, neighbor version 17
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  3 accepted prefixes
  2 announced prefixes


For address family: VPNv4 Flow Specification  BGP table version 5, neighbor version 5
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  1 announced prefixes

 Connections established 1; dropped 0
Local host: 1.1.1.1, Local port: 44219
Foreign host: 3.3.3.3, Foreign port: 179
Nexthop: 1.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 10.1.1.1, remote AS 100, local AS 65000, external link
  BGP version 4, local router ID 1.1.1.1, remote router ID 4.4.4.4
```

```
  BGP state = Established, up for 02:26:21
  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
    Address family IPv4 Flow Specification: advertised and received
  Received 372 messages, 0 notifications, 0 in queue
  Sent 359 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 6, neighbor version 6
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  0 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 17, neighbor version 17
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  4 announced prefixes

 Connections established 1; dropped 0
Local host: 10.1.1.2, Local port: 59044
Foreign host: 10.1.1.1, Foreign port: 179
Nexthop: 10.1.1.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 66.1.1.1, vrf vrf100, remote AS 100, local AS 65000, external link
  BGP version 4, local router ID 66.1.1.2, remote router ID 4.4.4.4
  BGP state = Established, up for 02:26:21
  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
    Address family IPv4 Flow Specification: advertised and received
  Received 356 messages, 0 notifications, 0 in queue
  Sent 362 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 0, Offset 0, Mask 0x1
  Community attribute sent to this neighbor (standard)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  5 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 1, neighbor version 1
  Index 0, Offset 0, Mask 0x1
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 66.1.1.2, Local port: 54954
Foreign host: 66.1.1.1, Foreign port: 179
Nexthop: 66.1.1.2
Nexthop global: ::
Nexthop local: ::
```

```
BGP connection: non shared network

PE1#
```

P1

The P1 node configurations are verified using the show commands listed below.

```
P1#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
PE1            xe7         b86a.9725.a7c2  Up     5         L1   IS-IS
                                           Up     5         L2   IS-IS
PE2            ce43        5c07.5819.486b  Up     8         L1   IS-IS
                                           Up     8         L2   IS-IS
P1#

P1#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
PE1.00-00             0x00000013   0x9F2C        680               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     PE1
  IP Address:   20.1.1.1
  Router ID:    1.1.1.1
  Metric:   10         IS-Extended PE1.02
    IPv4 Interface Address: 20.1.1.1
    Neighbor IP Address: 20.1.1.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.03
    IPv4 Interface Address: 40.1.1.1
    Neighbor IP Address: 40.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 20.1.1.0/24
  Metric:   0          IP-Extended 1.1.1.1/32
  Metric:   10         IP-Extended 40.1.1.0/24
```

© 2026 IP Infusion Inc. Proprietary

```
PE1.02-00             0x0000000F   0x6E4D        1141              0/0/0
  Metric:   0          IS-Extended PE1.00
  Metric:   0          IS-Extended P1.00
PE1.03-00             0x00000003   0xCFF2        344               0/0/0
  Metric:   0          IS-Extended PE1.00
  Metric:   0          IS-Extended P2.00
P1.00-00            * 0x00000013   0x6CF5        762               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     P1
  IP Address:   2.2.2.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended PE2.02
    IPv4 Interface Address: 11.1.1.2
    Neighbor IP Address: 11.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.02
    IPv4 Interface Address: 20.1.1.2
    Neighbor IP Address: 20.1.1.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 20.1.1.0/24
  Metric:   10         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 11.1.1.0/24
PE2.00-00             0x00000013   0x2F15        769               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     PE2
  IP Address:   11.1.1.1
  Router ID:    3.3.3.3
  Metric:   10         IS-Extended PE2.02
    IPv4 Interface Address: 11.1.1.1
    Neighbor IP Address: 11.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE2.01
    IPv4 Interface Address: 50.1.1.1
```

```
    Neighbor IP Address: 50.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   0          IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 11.1.1.0/24
PE2.01-00             0x0000000E   0xC9EB        1044              0/0/0
  Metric:   0          IS-Extended PE2.00
  Metric:   0          IS-Extended P2.00
PE2.02-00             0x0000000F   0x7047        1005              0/0/0
  Metric:   0          IS-Extended PE2.00
  Metric:   0          IS-Extended P1.00
P2.00-00              0x00000011   0x788C        749               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     P2
  IP Address:   6.6.6.6
  Router ID:    6.6.6.6
  Metric:   10         IS-Extended PE2.01
    IPv4 Interface Address: 50.1.1.2
    Neighbor IP Address: 50.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.03
    IPv4 Interface Address: 40.1.1.2
    Neighbor IP Address: 40.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 6.6.6.6/32
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IP-Extended 50.1.1.0/24

IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
PE1.00-00             0x00000017   0x40F9        680               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     PE1
```

```
  IP Address:   20.1.1.1
  Router ID:    1.1.1.1
  Metric:   10         IS-Extended PE1.02
    IPv4 Interface Address: 20.1.1.1
    Neighbor IP Address: 20.1.1.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.03
    IPv4 Interface Address: 40.1.1.1
    Neighbor IP Address: 40.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 20.1.1.0/24
  Metric:   0          IP-Extended 1.1.1.1/32
  Metric:   20         IP-Extended 2.2.2.2/32
  Metric:   20         IP-Extended 11.1.1.0/24
  Metric:   20         IP-Extended 3.3.3.3/32
  Metric:   20         IP-Extended 50.1.1.0/24
  Metric:   20         IP-Extended 6.6.6.6/32
  Metric:   10         IP-Extended 40.1.1.0/24
PE1.02-00             0x0000000F   0x6E4D        1141              0/0/0
  Metric:   0          IS-Extended PE1.00
  Metric:   0          IS-Extended P1.00
PE1.03-00             0x00000003   0xCFF2        344               0/0/0
  Metric:   0          IS-Extended PE1.00
  Metric:   0          IS-Extended P2.00
P1.00-00            * 0x00000019   0xCCF2        762               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     P1
  IP Address:   2.2.2.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended PE2.02
    IPv4 Interface Address: 11.1.1.2
    Neighbor IP Address: 11.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.02
```

```
    IPv4 Interface Address: 20.1.1.2
    Neighbor IP Address: 20.1.1.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 20.1.1.0/24
  Metric:   10         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 11.1.1.0/24
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   20         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 1.1.1.1/32
  Metric:   30         IP-Extended 6.6.6.6/32
  Metric:   20         IP-Extended 40.1.1.0/24
PE2.00-00            0x00000016   0x1BA1        769               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     PE2
  IP Address:   11.1.1.1
  Router ID:    3.3.3.3
  Metric:   10         IS-Extended PE2.02
    IPv4 Interface Address: 11.1.1.1
    Neighbor IP Address: 11.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE2.01
    IPv4 Interface Address: 50.1.1.1
    Neighbor IP Address: 50.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   0          IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 11.1.1.0/24
  Metric:   20         IP-Extended 2.2.2.2/32
  Metric:   20         IP-Extended 20.1.1.0/24
  Metric:   20         IP-Extended 1.1.1.1/32
  Metric:   20         IP-Extended 6.6.6.6/32
  Metric:   20         IP-Extended 40.1.1.0/24
PE2.01-00            0x0000000E   0xC9EB        1044              0/0/0
  Metric:   0          IS-Extended PE2.00
```

```
  Metric:   0          IS-Extended P2.00
PE2.02-00             0x0000000F   0x7047        1004              0/0/0
  Metric:   0          IS-Extended PE2.00
  Metric:   0          IS-Extended P1.00
P2.00-00              0x00000014   0x3E72        749               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     P2
  IP Address:   6.6.6.6
  Router ID:    6.6.6.6
  Metric:   10         IS-Extended PE2.01
    IPv4 Interface Address: 50.1.1.2
    Neighbor IP Address: 50.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.03
    IPv4 Interface Address: 40.1.1.2
    Neighbor IP Address: 40.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 6.6.6.6/32
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 1.1.1.1/32
  Metric:   30         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   20         IP-Extended 11.1.1.0/24
  Metric:   20         IP-Extended 20.1.1.0/24

P1#
```

P2

The CE1 node configurations are verified using the show commands listed below.

```
P2#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
PE1            ce51        b86a.9725.a7ee  Up     7         L1   IS-IS
                                           Up     7         L2   IS-IS
```

```
PE2            ce53        5c07.5819.4878      Up     7         L1   IS-IS
                                               Up     7         L2   IS-IS
P2#
```

```
P2#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
PE1.00-00             0x00000013   0x9F2C        620               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     PE1
  IP Address:   20.1.1.1
  Router ID:    1.1.1.1
  Metric:   10         IS-Extended PE1.02
    IPv4 Interface Address: 20.1.1.1
    Neighbor IP Address: 20.1.1.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.03
    IPv4 Interface Address: 40.1.1.1
    Neighbor IP Address: 40.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 20.1.1.0/24
  Metric:   0          IP-Extended 1.1.1.1/32
  Metric:   10         IP-Extended 40.1.1.0/24
PE1.02-00             0x0000000F   0x6E4D        1082              0/0/0
  Metric:   0          IS-Extended PE1.00
  Metric:   0          IS-Extended P1.00
PE1.03-00             0x00000004   0xCDF3        1185              0/0/0
  Metric:   0          IS-Extended PE1.00
  Metric:   0          IS-Extended P2.00
P1.00-00              0x00000013   0x6CF5        701               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     P1
  IP Address:   2.2.2.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended PE2.02
    IPv4 Interface Address: 11.1.1.2
    Neighbor IP Address: 11.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
```

```
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.02
    IPv4 Interface Address: 20.1.1.2
    Neighbor IP Address: 20.1.1.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 20.1.1.0/24
  Metric:   10         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 11.1.1.0/24
PE2.00-00            0x00000013   0x2F15        709              0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     PE2
  IP Address:   11.1.1.1
  Router ID:    3.3.3.3
  Metric:   10         IS-Extended PE2.02
    IPv4 Interface Address: 11.1.1.1
    Neighbor IP Address: 11.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE2.01
    IPv4 Interface Address: 50.1.1.1
    Neighbor IP Address: 50.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   0          IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 11.1.1.0/24
PE2.01-00            0x0000000E   0xC9EB        984              0/0/0
  Metric:   0          IS-Extended PE2.00
  Metric:   0          IS-Extended P2.00
```

```
PE2.02-00             0x0000000F   0x7047        945               0/0/0
  Metric:   0          IS-Extended PE2.00
  Metric:   0          IS-Extended P1.00
P2.00-00            * 0x00000011   0x788C        690               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     P2
  IP Address:   6.6.6.6
  Router ID:    6.6.6.6
  Metric:   10         IS-Extended PE2.01
    IPv4 Interface Address: 50.1.1.2
    Neighbor IP Address: 50.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.03
    IPv4 Interface Address: 40.1.1.2
    Neighbor IP Address: 40.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 6.6.6.6/32
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IP-Extended 50.1.1.0/24

IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
PE1.00-00             0x00000017   0x40F9        620               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     PE1
  IP Address:   20.1.1.1
  Router ID:    1.1.1.1
  Metric:   10         IS-Extended PE1.02
    IPv4 Interface Address: 20.1.1.1
    Neighbor IP Address: 20.1.1.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.03
    IPv4 Interface Address: 40.1.1.1
```

```
    Neighbor IP Address: 40.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 20.1.1.0/24
  Metric:   0          IP-Extended 1.1.1.1/32
  Metric:   20         IP-Extended 2.2.2.2/32
  Metric:   20         IP-Extended 11.1.1.0/24
  Metric:   20         IP-Extended 3.3.3.3/32
  Metric:   20         IP-Extended 50.1.1.0/24
  Metric:   20         IP-Extended 6.6.6.6/32
  Metric:   10         IP-Extended 40.1.1.0/24
PE1.02-00             0x0000000F   0x6E4D        1082              0/0/0
  Metric:   0          IS-Extended PE1.00
  Metric:   0          IS-Extended P1.00
PE1.03-00             0x00000004   0xCDF3        1185              0/0/0
  Metric:   0          IS-Extended PE1.00
  Metric:   0          IS-Extended P2.00
P1.00-00              0x00000019   0xCCF2        701               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     P1
  IP Address:   2.2.2.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended PE2.02
    IPv4 Interface Address: 11.1.1.2
    Neighbor IP Address: 11.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.02
    IPv4 Interface Address: 20.1.1.2
    Neighbor IP Address: 20.1.1.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 20.1.1.0/24
  Metric:   10         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 11.1.1.0/24
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   20         IP-Extended 50.1.1.0/24
```

```
  Metric:   10         IP-Extended 1.1.1.1/32
  Metric:   30         IP-Extended 6.6.6.6/32
  Metric:   20         IP-Extended 40.1.1.0/24
PE2.00-00            0x00000016   0x1BA1        710               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     PE2
  IP Address:   11.1.1.1
  Router ID:    3.3.3.3
  Metric:   10         IS-Extended PE2.02
    IPv4 Interface Address: 11.1.1.1
    Neighbor IP Address: 11.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE2.01
    IPv4 Interface Address: 50.1.1.1
    Neighbor IP Address: 50.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   0          IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 11.1.1.0/24
  Metric:   20         IP-Extended 2.2.2.2/32
  Metric:   20         IP-Extended 20.1.1.0/24
  Metric:   20         IP-Extended 1.1.1.1/32
  Metric:   20         IP-Extended 6.6.6.6/32
  Metric:   20         IP-Extended 40.1.1.0/24
PE2.01-00            0x0000000E   0xC9EB        984               0/0/0
  Metric:   0          IS-Extended PE2.00
  Metric:   0          IS-Extended P2.00
PE2.02-00            0x0000000F   0x7047        945               0/0/0
  Metric:   0          IS-Extended PE2.00
  Metric:   0          IS-Extended P1.00
P2.00-00           * 0x00000014   0x3E72        690               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     P2
  IP Address:   6.6.6.6
  Router ID:    6.6.6.6
  Metric:   10         IS-Extended PE2.01
    IPv4 Interface Address: 50.1.1.2
    Neighbor IP Address: 50.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
```

```
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.03
    IPv4 Interface Address: 40.1.1.2
    Neighbor IP Address: 40.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 6.6.6.6/32
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 1.1.1.1/32
  Metric:   30         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   20         IP-Extended 11.1.1.0/24
  Metric:   20         IP-Extended 20.1.1.0/24

P2#
```

PE2

The PE2 node configurations are verified using the show commands listed below.

```
PE2#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
P1             ce43        3417.ebe4.af2c  Up     22        L1   IS-IS
                                           Up     22        L2   IS-IS
P2             ce4         3c2c.99c7.07af  Up     21        L1   IS-IS
                                           Up     21        L2   IS-IS
PE2#

PE2#sh ip bgp ipv4 flowspec summary
BGP router identifier 3.3.3.3, local AS number 65000
BGP table version is 18
3 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 65000  394        409      18      0      0  02:35:22              2
30.1.1.1                4   200  382        372      18      0      0  02:22:01              2

Total number of neighbors 2

Total number of Established sessions 2
BGP router identifier 55.1.1.1, local AS number 65000
BGP VRF vrf100 Route Distinguisher: 100:1
BGP table version is 1
```

```
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
55.1.1.2                 4   200  365        373       1      0      0  02:22:01              2

Total number of neighbors 1

Total number of Established sessions 1
PE2#
PE2#
PE2#show ip bgp ipv4 flowspec
BGP table version is 18, local router ID is 3.3.3.3
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
*>i  flowspec/136     1.1.1.1               0        100       0   100 i
*>   flowspec/200     30.1.1.1              0        100       0   200 i
*>   flowspec/152     0.0.0.0               0        100       32768  i
*>i  flowspec/48      1.1.1.1               0        100       0    i
*>   flowspec/152     30.1.1.1              0        100       0   200 i

Total number of prefixes 5
PE2#
PE2#show ip bgp ipv4 flowspec detail
BGP table version is 18, local router ID is 3.3.3.3
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
*>i  flowspec/136     1.1.1.1               0        100       0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>   flowspec/200     30.1.1.1              0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*>   flowspec/152     0.0.0.0               0        100       32768  i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
```

```
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*>i  flowspec/48      1.1.1.1              0        100       0    i

  FlowSpec:
    destination ip address: 4.4.4.4/32

  FlowSpec Actions:
    rt_redirect: as_type: as 100:10 vrf_id:0 fib_id:0

*>   flowspec/152     30.1.1.1             0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 5
PE2#
PE2#
PE2#sh ip bgp ipv4 flowspec vrf all
BGP table version is 1, local router ID is 55.1.1.1
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

BGP Route Table for VRF vrf100
*>i  flowspec/136     1.1.1.1              0        100       0   100 i
*>   flowspec/200     55.1.1.2             0        100       0   200 i
*>   flowspec/152     55.1.1.2             0        100       0   200 i

Total number of prefixes 3
*>i  flowspec/136     1.1.1.1              0        100       0   100 i
*>   flowspec/200     30.1.1.1             0        100       0   200 i
*>   flowspec/152     0.0.0.0              0        100       32768  i
*>i  flowspec/48      1.1.1.1              0        100       0    i
*>   flowspec/152     30.1.1.1             0        100       0   200 i

Total number of prefixes 5
PE2#
PE2#
PE2#sh ip bgp ipv4 flowspec vrf all detail
BGP table version is 1, local router ID is 55.1.1.1
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

BGP Route Table for VRF vrf100
*>i  flowspec/136     1.1.1.1              0        100       0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
```

```
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>   flowspec/200     55.1.1.2             0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*>   flowspec/152     55.1.1.2             0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 3
*>i  flowspec/136     1.1.1.1              0        100       0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>   flowspec/200     30.1.1.1             0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e
```

```
*>   flowspec/152     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*>i  flowspec/48      1.1.1.1              0        100       0    i

  FlowSpec:
    destination ip address: 4.4.4.4/32

  FlowSpec Actions:
    rt_redirect: as_type: as 100:10 vrf_id:0 fib_id:0

*>   flowspec/152     30.1.1.1             0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 5
PE2#
PE2#sh ip bgp vpnv4 all
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
Route Distinguisher: 100:1 (Default for VRF vrf100)
*>i  4.4.4.4/32       1.1.1.1              0        100       0   100 ?
*> l 5.5.5.5/32       55.1.1.2             0        100       0   200 ?
*>i  10.1.1.0/24      1.1.1.1              0        100       0   100 ?
*> l 30.1.1.0/24      55.1.1.2             0        100       0   200 ?
*> l 55.1.1.0/24      0.0.0.0              0        100       32768  ?
*                     55.1.1.2             0        100       0   200 ?
*>i  66.1.1.0/24      1.1.1.1              0        100       0    ?
*>i  100.1.1.0/24     1.1.1.1              0        100       0   100 ?
*> l 200.1.1.0        55.1.1.2             0        100       0   200 ?
 Announced routes count = 5
 Accepted routes count = 4
Route Distinguisher: 100:1
*>i  4.4.4.4/32       1.1.1.1              0        100       0   100 ?
*>i  10.1.1.0/24      1.1.1.1              0        100       0   100 ?
*>i  66.1.1.0/24      1.1.1.1              0        100       0    ?
*>i  100.1.1.0/24     1.1.1.1              0        100       0   100 ?
 Announced routes count = 0
 Accepted routes count = 4
PE2#
PE2#sh ip bgp vpnv4 flowspec summary
BGP router identifier 3.3.3.3, local AS number 65000
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
1.1.1.1                  4 65000  397        413       5      0      0  02:36:40              1

Total number of neighbors 1

Total number of Established sessions 1
PE2#
PE2#
PE2#sh ip bgp vpnv4 flowspec vrf vrf100
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
Route Distinguisher: 100:1 (Default for VRF vrf100)
*>i  flowspec/136     1.1.1.1              0        100        0   100 i
*>   flowspec/200     55.1.1.2             0        100        0   200 i
*>   flowspec/152     55.1.1.2             0        100        0   200 i
 Announced routes count = 2
 Accepted routes count = 1
PE2#
PE2#sh ip bgp vpnv4 flowspec vrf vrf100 detail
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
Route Distinguisher: 100:1 (Default for VRF vrf100)
*>i  flowspec/136     1.1.1.1              0        100        0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>   flowspec/200     55.1.1.2             0        100        0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*>   flowspec/152     55.1.1.2             0        100        0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
```

```
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400

  FlowSpec Actions:
    traffic_rate: 100000 bps

 Announced routes count = 2
 Accepted routes count = 1
PE2#
```

```
PE2#show ip bgp neighbors
BGP neighbor is 1.1.1.1, remote AS 65000, local AS 65000, internal link
  BGP version 4, local router ID 3.3.3.3, remote router ID 1.1.1.1
  BGP state = Established, up for 02:34:56
  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
    Address family VPNv4 Unicast: advertised and received
    Address family IPv4 Flow Specification: advertised and received
    Address family VPNv4 Flow Specification: advertised and received
  Received 392 messages, 1 notifications, 0 in queue
  Sent 408 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 7, neighbor version 7
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: VPNv4 Unicast  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 18, neighbor version 18
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  3 announced prefixes


For address family: VPNv4 Flow Specification  BGP table version 5, neighbor version 5
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  1 accepted prefixes
  2 announced prefixes

 Connections established 2; dropped 1
Local host: 3.3.3.3, Local port: 179
Foreign host: 1.1.1.1, Foreign port: 44219
Nexthop: 3.3.3.3
Nexthop global: ::
```

```
Nexthop local: ::
BGP connection: non shared network
Last Reset: 02:34:58, due to BGP Notification received
Notification Error Message: (Cease/Peer Unconfigured.)

BGP neighbor is 30.1.1.1, remote AS 200, local AS 65000, external link
  BGP version 4, local router ID 3.3.3.3, remote router ID 5.5.5.5
  BGP state = Established, up for 02:21:35
  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
    Address family IPv4 Flow Specification: advertised and received
  Received 381 messages, 0 notifications, 0 in queue
  Sent 369 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 7, neighbor version 7
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  0 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 18, neighbor version 18
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  3 announced prefixes

 Connections established 2; dropped 1
Local host: 30.1.1.2, Local port: 179
Foreign host: 30.1.1.1, Foreign port: 58782
Nexthop: 30.1.1.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 02:21:35, due to OPEN Message Error (Notification sent)
Notification Error Message: (OPEN Message Error/Bad BGP Identifier. Reason: )

BGP neighbor is 55.1.1.2, vrf vrf100, remote AS 200, local AS 65000, external link
  BGP version 4, local router ID 55.1.1.1, remote router ID 5.5.5.5
  BGP state = Established, up for 02:21:35
  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
    Address family IPv4 Flow Specification: advertised and received
  Received 364 messages, 0 notifications, 0 in queue
  Sent 370 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 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (standard)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  5 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
```

```
  2 accepted prefixes
  1 announced prefixes

 Connections established 2; dropped 1
Local host: 55.1.1.1, Local port: 179
Foreign host: 55.1.1.2, Foreign port: 41566
Nexthop: 55.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 02:21:35, due to OPEN Message Error (Notification sent)
Notification Error Message: (OPEN Message Error/Bad BGP Identifier. Reason: )

PE2#
PE2#
```

```
PE2#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
PE1.00-00             0x00000013   0x9F2C        529               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     PE1
  IP Address:   20.1.1.1
  Router ID:    1.1.1.1
  Metric:   10         IS-Extended PE1.02
    IPv4 Interface Address: 20.1.1.1
    Neighbor IP Address: 20.1.1.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.03
    IPv4 Interface Address: 40.1.1.1
    Neighbor IP Address: 40.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 20.1.1.0/24
  Metric:   0          IP-Extended 1.1.1.1/32
  Metric:   10         IP-Extended 40.1.1.0/24
PE1.02-00             0x0000000F   0x6E4D        990               0/0/0
  Metric:   0          IS-Extended PE1.00
  Metric:   0          IS-Extended P1.00
PE1.03-00             0x00000004   0xCDF3        1093              0/0/0
  Metric:   0          IS-Extended PE1.00
  Metric:   0          IS-Extended P2.00
P1.00-00              0x00000013   0x6CF5        609               0/0/0
  Area Address: 49
```

```
  NLPID:        0xCC
  Hostname:     P1
  IP Address:   2.2.2.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended PE2.02
    IPv4 Interface Address: 11.1.1.2
    Neighbor IP Address: 11.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.02
    IPv4 Interface Address: 20.1.1.2
    Neighbor IP Address: 20.1.1.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 20.1.1.0/24
  Metric:   10         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 11.1.1.0/24
PE2.00-00           * 0x00000013   0x2F15        619               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     PE2
  IP Address:   11.1.1.1
  Router ID:    3.3.3.3
  Metric:   10         IS-Extended PE2.02
    IPv4 Interface Address: 11.1.1.1
    Neighbor IP Address: 11.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE2.01
    IPv4 Interface Address: 50.1.1.1
    Neighbor IP Address: 50.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
```

```
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   0          IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 11.1.1.0/24
PE2.01-00           * 0x0000000E   0xC9EB        893               0/0/0
  Metric:   0          IS-Extended PE2.00
  Metric:   0          IS-Extended P2.00
PE2.02-00           * 0x0000000F   0x7047        854               0/0/0
  Metric:   0          IS-Extended PE2.00
  Metric:   0          IS-Extended P1.00
P2.00-00             0x00000011   0x788C        597               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     P2
  IP Address:   6.6.6.6
  Router ID:    6.6.6.6
  Metric:   10         IS-Extended PE2.01
    IPv4 Interface Address: 50.1.1.2
    Neighbor IP Address: 50.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.03
    IPv4 Interface Address: 40.1.1.2
    Neighbor IP Address: 40.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 6.6.6.6/32
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IP-Extended 50.1.1.0/24

IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
PE1.00-00             0x00000017   0x40F9        529               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     PE1
  IP Address:   20.1.1.1
  Router ID:    1.1.1.1
  Metric:   10         IS-Extended PE1.02
    IPv4 Interface Address: 20.1.1.1
    Neighbor IP Address: 20.1.1.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
```

```
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.03
    IPv4 Interface Address: 40.1.1.1
    Neighbor IP Address: 40.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 20.1.1.0/24
  Metric:   0          IP-Extended 1.1.1.1/32
  Metric:   20         IP-Extended 2.2.2.2/32
  Metric:   20         IP-Extended 11.1.1.0/24
  Metric:   20         IP-Extended 3.3.3.3/32
  Metric:   20         IP-Extended 50.1.1.0/24
  Metric:   20         IP-Extended 6.6.6.6/32
  Metric:   10         IP-Extended 40.1.1.0/24
PE1.02-00             0x0000000F   0x6E4D        990              0/0/0
  Metric:   0          IS-Extended PE1.00
  Metric:   0          IS-Extended P1.00
PE1.03-00             0x00000004   0xCDF3        1093             0/0/0
  Metric:   0          IS-Extended PE1.00
  Metric:   0          IS-Extended P2.00
P1.00-00              0x00000019   0xCCF2        609              0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     P1
  IP Address:   2.2.2.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended PE2.02
    IPv4 Interface Address: 11.1.1.2
    Neighbor IP Address: 11.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.02
    IPv4 Interface Address: 20.1.1.2
    Neighbor IP Address: 20.1.1.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
```

```
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 20.1.1.0/24
  Metric:   10         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 11.1.1.0/24
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   20         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 1.1.1.1/32
  Metric:   30         IP-Extended 6.6.6.6/32
  Metric:   20         IP-Extended 40.1.1.0/24
PE2.00-00           * 0x00000016   0x1BA1        619               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     PE2
  IP Address:   11.1.1.1
  Router ID:    3.3.3.3
  Metric:   10         IS-Extended PE2.02
    IPv4 Interface Address: 11.1.1.1
    Neighbor IP Address: 11.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE2.01
    IPv4 Interface Address: 50.1.1.1
    Neighbor IP Address: 50.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   0          IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 11.1.1.0/24
  Metric:   20         IP-Extended 2.2.2.2/32
  Metric:   20         IP-Extended 20.1.1.0/24
  Metric:   20         IP-Extended 1.1.1.1/32
  Metric:   20         IP-Extended 6.6.6.6/32
  Metric:   20         IP-Extended 40.1.1.0/24
PE2.01-00           * 0x0000000E   0xC9EB        893               0/0/0
  Metric:   0          IS-Extended PE2.00
  Metric:   0          IS-Extended P2.00
PE2.02-00           * 0x0000000F   0x7047        854               0/0/0
  Metric:   0          IS-Extended PE2.00
  Metric:   0          IS-Extended P1.00
P2.00-00              0x00000014   0x3E72        598               0/0/0
  Area Address: 49
  NLPID:        0xCC
  Hostname:     P2
```

```
  IP Address:   6.6.6.6
  Router ID:    6.6.6.6
  Metric:   10         IS-Extended PE2.01
    IPv4 Interface Address: 50.1.1.2
    Neighbor IP Address: 50.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended PE1.03
    IPv4 Interface Address: 40.1.1.2
    Neighbor IP Address: 40.1.1.1
    Maximum Link Bandwidth : 100000000.00 kbits/sec
    Reservable Bandwidth : 100000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 6.6.6.6/32
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 1.1.1.1/32
  Metric:   30         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   20         IP-Extended 11.1.1.0/24
  Metric:   20         IP-Extended 20.1.1.0/24

PE2#
```

CE2

The CE2 node configurations are verified using the show commands listed below.

```
CE2#show ip bgp
BGP table version is 2, local router ID is 5.5.5.5
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
*>   4.4.4.4/32       55.1.1.1              0        100       0   65000 100 ?
*>   5.5.5.5/32       0.0.0.0               0        100       32768  ?
*>   10.1.1.0/24      55.1.1.1              0        100       0   65000 100 ?
*>   30.1.1.0/24      0.0.0.0               0        100       32768  ?
*>   55.1.1.0/24      0.0.0.0               0        100       32768  ?
*                     55.1.1.1              0        100       0   65000 ?
*>   66.1.1.0/24      55.1.1.1              0        100       0   65000 ?
*>   100.1.1.0/24     55.1.1.1              0        100       0   65000 100 ?
*>   200.1.1.0        0.0.0.0               0        100       32768  ?

Total number of prefixes 8
```

```
CE2#

CE2#sh ip bgp ipv4 flowspec summary
BGP router identifier 5.5.5.5, local AS number 200
BGP table version is 7
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
30.1.1.2                 4 65000  356        365       7      0      0  02:25:26               3
55.1.1.1                 4 65000  355        352       7      0      0  02:25:26               1

Total number of neighbors 2

Total number of Established sessions 2
CE2#sh ip bgp ipv4 flowspec
BGP table version is 7, local router ID is 5.5.5.5
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
*>   flowspec/136     55.1.1.1              0        100       0   65000 100 i
*                     30.1.1.2              0        100       0   65000 100 i
*>   flowspec/200     0.0.0.0               0        100       32768  i
*    flowspec/152     30.1.1.2              0        100       0   65000 i
*    flowspec/48      30.1.1.2              0        100       0   65000 i
*>   flowspec/152     0.0.0.0               0        100       32768  i

Total number of prefixes 5
CE2#
CE2#sh ip bgp ipv4 flowspec detail
BGP table version is 7, local router ID is 5.5.5.5
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
*>   flowspec/136     55.1.1.1              0        100       0   65000 100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*                     30.1.1.2              0        100       0   65000 100 i
  FlowSpec Actions:
    traffic_rate: 21000 bps

*>   flowspec/200     0.0.0.0               0        100       32768  i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21
```

© 2026 IP Infusion Inc. Proprietary

```
  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*    flowspec/152     30.1.1.2             0        100       0   65000 i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*    flowspec/48      30.1.1.2             0        100       0   65000 i

  FlowSpec:
    destination ip address: 4.4.4.4/32

  FlowSpec Actions:
    rt_redirect: as_type: as 100:10 vrf_id:0 fib_id:0

*>   flowspec/152     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 5
CE2#
CE2#
CE2#sh ip bgp ipv4 flowspec vrf all
BGP table version is 7, local router ID is 5.5.5.5
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
*>   flowspec/136     55.1.1.1              0        100       0   65000 100 i
*                     30.1.1.2              0        100       0   65000 100 i
*>   flowspec/200     0.0.0.0               0        100       32768  i
*    flowspec/152     30.1.1.2              0        100       0   65000 i
*    flowspec/48      30.1.1.2              0        100       0   65000 i
*>   flowspec/152     0.0.0.0               0        100       32768  i

Total number of prefixes 5
CE2#
CE2#
CE2#sh ip bgp ipv4 flowspec vrf all detail
BGP table version is 7, local router ID is 5.5.5.5
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
*>   flowspec/136     55.1.1.1              0        100       0   65000 100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
```

```
    source ip address: 10.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*                     30.1.1.2             0        100       0   65000 100 i
  FlowSpec Actions:
    traffic_rate: 21000 bps

*>   flowspec/200     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*    flowspec/152     30.1.1.2             0        100       0   65000 i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*    flowspec/48      30.1.1.2             0        100       0   65000 i

  FlowSpec:
    destination ip address: 4.4.4.4/32

  FlowSpec Actions:
    rt_redirect: as_type: as 100:10 vrf_id:0 fib_id:0

*>   flowspec/152     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 5
CE2#
```

```
CE2#sh ip bgp neighbors
BGP neighbor is 30.1.1.2, remote AS 65000, local AS 200, external link
  BGP version 4, local router ID 5.5.5.5, remote router ID 3.3.3.3
  BGP state = Established, up for 02:25:18
  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
    Address family IPv4 Flow Specification: advertised and received
  Received 355 messages, 1 notifications, 0 in queue
  Sent 365 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)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  8 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 7, neighbor version 7
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  3 accepted prefixes
  3 announced prefixes

 Connections established 1; dropped 0
Local host: 30.1.1.1, Local port: 58782
Foreign host: 30.1.1.2, Foreign port: 179
Nexthop: 30.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 02:25:23, due to BGP Notification received
Notification Error Message: (OPEN Message Error/Bad BGP Identifier. Reason: )

BGP neighbor is 55.1.1.1, remote AS 65000, local AS 200, external link
  BGP version 4, local router ID 5.5.5.5, remote router ID 55.1.1.1
  BGP state = Established, up for 02:25:18
  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
    Address family IPv4 Flow Specification: advertised and received
  Received 354 messages, 1 notifications, 0 in queue
  Sent 351 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 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  5 accepted prefixes
  4 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 7, neighbor version 7
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  1 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
```

```
Local host: 55.1.1.2, Local port: 41566
Foreign host: 55.1.1.1, Foreign port: 179
Nexthop: 55.1.1.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 02:25:23, due to BGP Notification received
Notification Error Message: (OPEN Message Error/Bad BGP Identifier. Reason: )

CE2#
CE2#
```

BGP Flowspec Configuration With RR

Topology

Figure 89. BGP Flowpec support for IPv4 with RR

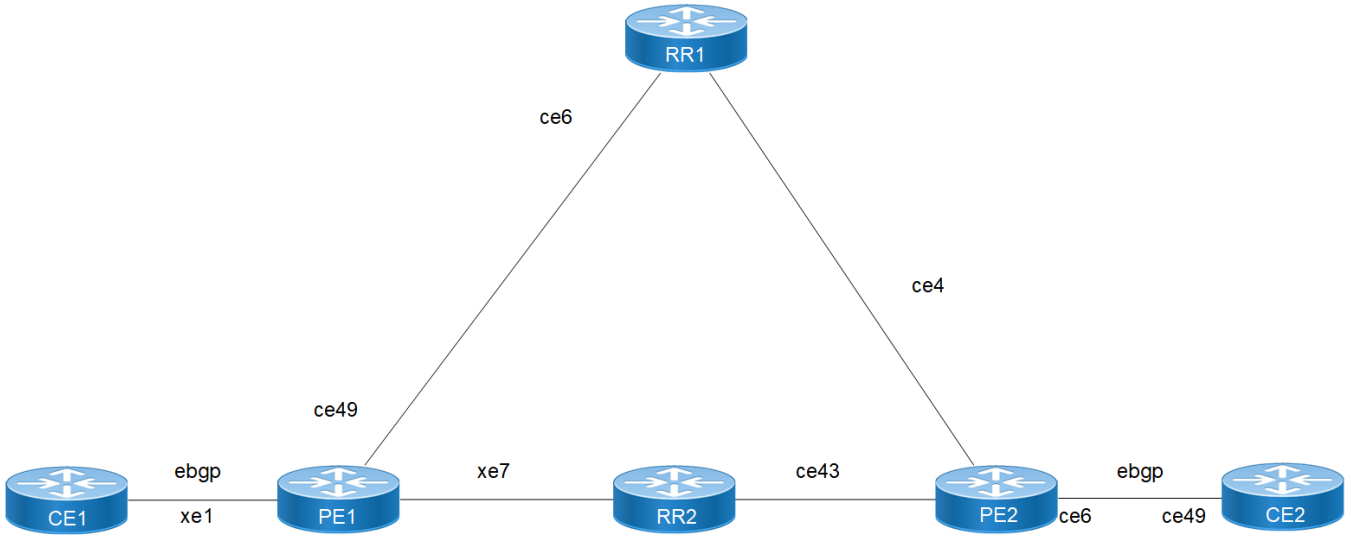

CE1 Configuration

`CE1#configure terminal`	Enter configure mode.
`CE1(config)# class-map type traffic match-all ocnos`	Configure class map with class name
`CE1(config-cmap-tr)#match destination-address ipv4 5.5.5.5/32`	Configure match with destination ip address
`CE1(config-cmap-tr)#match source-address ipv4 4.4.4.4/32`	Enable match with source ip address
`CE1(config-cmap-tr)#match protocol 12`	Enable match with protocol
`CE1(config-cmap-tr)#match port 333-444`	Enable match with port
`CE1(config-cmap-tr)#exit`	Exit from class map traffic
`CE1(config)#commit`	Commits the configuration
`CE1(config)#class-map type traffic match-all ocnos2`	Configure class map with class name
`CE1(config-cmap-tr)# match destination-address ipv4`	Configure match with destination ip address

`10.1.1.0/24`	
`CE1(config-cmap-tr)# match source-address ipv4 10.1.1.0/24`	Enable match with source ip address
`CE1(config-cmap-tr)# match protocol 44`	Enable match with protocol
`CE1(config-cmap-tr)# match destination-port 44-144`	Enable match with destination port
`CE1(config-cmap-tr)# match source-port 22-122`	Enter match with source port
`CE1(config-cmap-tr)#exit`	Exit from class map traffic
`CE1(config)#commit`	Commits the configuration
`CE1(config)#policy-map type pbr CE1`	Enable policy map type pbr policy name
`CE1(config-pmap-pbr)# class type traffic ocnos2`	Enable class type traffic class name
`CE1(config-pmap-pbr-c)# police rate 21000 bps`	Set the policy rate
`CE1(config-pmap-pbr-c)#exit`	Exit from class
`CE1(config-pmap-pbr)#exit`	Exit from policy map
`CE1(config)#policy-map type pbr PE1`	Enable policy map type pbr policy name
`CE1(config-pmap-pbr)# class type traffic ocnos`	Enable class type traffic class name
`CE1(config-pmap-pbr-c)# police rate 4444444 bps`	Set the policy rate
`CE1(config-cmap-tr)#exit`	Exit from class map traffic
`CE1(config)#commit`	Commits the configuration
`CE1(config)#flowspec`	Configure flowspec
`CE1(config-flowspec)#address-family ipv4`	Enter address family ipv4 mode
`CE1(config-flowspec-af)#service-policy type pbr PE1`	Enable service policy type pbr with policy name
`CE1(config-flowspec-af)#service-policy type pbr CE1`	Enable service policy type pbr with policy name
`CE1(config-flowspec-af)#exit`	Exit from address family
`CE1(config-flowspec)#exit`	Exit from flowspec mode
`CE1(config)#commit`	Commits the configuration
`CE1(config)#interface lo`	Enter the loopback onterface mode
`CE1(config-if)# ip address 4.4.4.4/32 secondary`	Assign the ip add as secondary
`CE1(config-if)# ip router isis 1`	Enable the router isis instance
`CE1(config-if)#exit`	Exit from interface mode
`CE1(config)#commit`	Commits the configuration
`CE1(config)#interface xe1`	Enter the interface mode
`CE1(config-if)#ip address 10.1.1.1/24`	Configure ip address
`CE1(config)#interface xe1.10`	Configure subinterface
`CE1(config-if)# encapsulation dot1q 10`	Enable encapsulate dot1q vlan id
`CE1(config-if)# ip address 66.1.1.1/24`	Configure ip address
`CE1(config-if)#exit`	Exit from interface mode

© 2026 IP Infusion Inc. Proprietary

`CE1(config)#commit`	Commits the configuration
`CE1(config)#router bgp 100`	Configure bgp intstance
`CE1(config-router)# neighbor 10.1.1.2 remote-as 65000`	Enable neighbour ip address with remote-as 65000
`CE1(config-router)# neighbor 66.1.1.2 remote-as 65000`	Enable neighbour ip address with remote-as 65000
`CE1(config-router)# address-family ipv4 unicast`	Enter the address family ipv4 unicast
`CE1(config-router-af)# redistribute connected`	Enter redistribute connected
`CE1(config-router-af)# neighbor 10.1.1.2 activate`	Activate the neighbour
`CE1(config-router-af)# neighbor 66.1.1.2 activate`	Activate the neighbour
`CE1(config-router-af)# exit-address-family`	Exit from address family mode
`CE1(config-router)# address-family ipv4 flowspec`	Enter the address family ipv4 flowspec
`CE1(config-router-af)# neighbor 10.1.1.2 activate`	Activate the neighbour
`CE1(config-router-af)#neighbor 66.1.1.2 activate`	Activate the neighbour
`CE1(config-router-af)# exit-address-family`	Exit from address family mode
`CE1(config-router)#exit`	Exit from router mode
`CE1(config)#commit`	Commits the configuration

PE1 Configuration

`PE1#configure terminal`	Enter configure mode.
`PE1(config)#ip vrf vrf100`	Enter vrf vrf name
`PE1(config-vrf)# rd 100:1`	Enter rd value
`PE1(config-vrf)# route-target both 100:1`	Enter router targeted as both
`PE1(config-vrf)#exit`	Exit from vrf mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#ip vrf vrf101`	Enter vrf vrf name
`PE1(config-vrf)# rd 101:1`	Enter rd value
`PE1(config-vrf)# route-target both 101:1`	Enter router targeted as both
`PE1(config-vrf)#exit`	Exit from vrf mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#qos enable`	Enable qos
`PE1(config)#qos statistics`	Enable qos statistics
`PE1-7026(config)#class-map type traffic match-all class1`	Configure class map with class name
`PE1-7026(config-cmap-tr)# match destination-address ipv4 3.3.3.3/32`	Configure match with destination ip address
`PE1-7026(config-cmap-tr)# match source-address ipv4 1.1.1.1/32`	Enable match with source ip address
`PE1-7026(config-cmap-tr)# match protocol 78`	Enable match with protocol

© 2026 IP Infusion Inc. Proprietary

`PE1-7026(config-cmap-tr)# match port 101-201`	Enable match with port range
`PE1-7026(config-cmap-tr)# match destination-port 400-499`	Enable match with destination-port range
`PE1-7026(config-cmap-tr)# match source-port 330-399`	Enable match with source-port range
`PE1-7026(config-cmap-tr)# match ipv4 icmp-code 10`	Enable match with icmp code
`PE1-7026(config-cmap-tr)# match packet-length 22-222`	Enable match with packet length range
`PE1-7026(config-cmap-tr)# match dscp 10`	Enable match with dscp
`PE1-7026(config-cmap-tr)#exit`	Exit from class map traffic
`PE1-7026(config)#commit`	Commits the configuration
`PE1-7026(config)#policy-map type pbr pol11`	Configure policy map type pbr with policy name
`PE1-7026(config-pmap-pbr)# class type traffic class1`	Enable class type traffic with class map name
`PE1-7026(config-pmap-pbr-c)# police rate 1000011 bps`	Enable policy rate
`PE1-7026(config-pmap-pbr-c)# set dscp 0e`	Set the dscp value
`PE1-7026(config-pmap-pbr-c)#exit`	Exit from polocy map traffic
`PE1-7026(config-pmap-pbr)#exit`	Exit from policy map
`PE1-7026(config)#commit`	Commits the configuration
`PE1(config)#flowspec`	Configure flowspec
`PE1(config-flowspec)# address-family ipv4`	Enter address-family mode with ipv4
`PE1(config-flowspec-af)# service-policy type pbr pol11`	Enable service plocy type pbr with policy name
`PE1(config-flowspec-af)#exit`	Exit from address family mode
`PE1(config-flowspec)#exit`	Exit from flowspec mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#router ldp`	Enter router ldp mode
`PE1(config-router)# targeted-peer ipv4 3.3.3.3`	Enable targeted peer ipv4 address
`PE1(config-router-targeted-peer)# exit-targeted-peer-mode`	Exit from targeted peer mode
`PE1(config-router)#exit`	Exit from router ldp
`PE1(config)#commit`	Commits the configuration
`PE1(config)#interface lo`	Enter interface mode
`PE1(config-if)# ip address 1.1.1.1/32 secondary`	Assign the ip address secondary
`PE1(config-if)# ip router isis 1`	Enable router isis instance
`PE1(config-if)#exit`	Exit from interface mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#interface ce49`	Enter interface mode
`PE1(config-if)# ip address 40.1.1.1/24`	Assign the ip add

© 2026 IP Infusion Inc. Proprietary

`PE1(config-if)# label-switching`	Enable label switching
`PE1(config-if)# ip router isis 1`	Enter router isis instance
`PE1(config-if)# enable-ldp ipv4`	Enable ldp ipv4
`PE1(config-if)#exit`	Exit from interface mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#interface xe1`	Enter the interface mode
`PE1(config-if)# ip address 10.1.1.2/24`	Assign the ip add
`PE1(config-if)#exit`	Exit from interface mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#interface xe1.10`	Enter the sub-interface mode
`PE1(config-if)# encapsulation dot1q 10`	Enable encapsulate dot1q with vlan id
`PE1(config-if)# ip vrf forwarding vrf100`	Enable ip vrf forwarding vrf name
`PE1(config-if)# ip address 66.1.1.2/24`	Assign the ip add
`PE1(config-if)#exit`	Exit from interface mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#interface xe7`	Enter interface mode
`PE1(config-if)# ip address 20.1.1.1/24`	Assign the ip add
`PE1(config-if)# label-switching`	Enable label switching
`PE1(config-if)# ip router isis 1`	Enter router isis instance
`PE1(config-if)# enable-ldp ipv4`	Enable ldp ipv4
`PE1(config-if)#exit`	Exit from interface mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#router isis 1`	Configure isis instance
`PE1(config-router)# is-type level-1-2`	Configure is-type-level-1-2
`PE1(config-router)# metric-style wide`	Configure metric style wide
`PE1(config-router)# mpls traffic-eng router-id 1.1.1.1`	Configure router id
`PE1(config-router)# mpls traffic-eng level-1`	Configure traffic engg level-1
`PE1(config-router)# mpls traffic-eng level-2`	Configure traffic engg level-2
`PE1(config-router)# dynamic-hostname`	Configure hostname
`PE1(config-router)# bfd all-interfaces`	Enable bfd all interfaces
`PE1(config-router)# net 49.0000.0000.0001.00`	Configure network id
`PE1 (config-router)# passive-interface lo`	Configure passive interface as loopback
`PE1(config-router)#exit`	Exit from isis instance mode
`PE1(config)#commit`	Commits the configuration
`PE1(config)#router bgp 65000`	Configure bgp instance
`PE1(config-router)# neighbor 6.6.6.6 remote-as`	Enable neighbour address with remote-as 65000

© 2026 IP Infusion Inc. Proprietary

65000	
PE1(config-router)# neighbor 10.1.1.1 remote-as 100	Enable neighbour address with remote-as 100
PE1(config-router)# neighbor 20.1.1.2 remote-as 65000	Enable neighbour address with remote-as 65000
PE1(config-router)# neighbor 6.6.6.6 update-source lo	Enable neighbour address with update loopback
PE1(config-router)# address-family ipv4 unicast	Enter address family ipv4 mode
PE1(config-router-af)# neighbor 6.6.6.6 activate	Activate the neighbour
PE1(config-router-af)# neighbor 6.6.6.6 next-hop-self	Enter neighbor ip add with next hop self
PE1(config-router-af)# neighbor 10.1.1.1 activate	Activate the neighbour
PE1(config-router-af)# neighbor 20.1.1.2 activate	Activate the neighbour
PE1(config-router-af)# neighbor 20.1.1.2 next-hop-self	Enter neighbor ip add with next hop self
PE1(config-router-af)# exit-address-family	Exit from address family
PE1(config-router)# address-family vpnv4 unicast	Enter address family vpnv4 unicast mode
PE1(config-router-af)# neighbor 6.6.6.6 activate	Activate the neighbour
PE1(config-router-af)# neighbor 6.6.6.6 next-hop-self	Enter neighbor ip add with next hop self
PE1(config-router-af)# neighbor 20.1.1.2 activate	Activate the neighbour
PE1(config-router-af)# neighbor 20.1.1.2 next-hop-self	Enter neighbor ip add with next hop self
PE1(config-router-af)# exit-address-family	Exit from address family
PE1(config-router)# address-family ipv4 flowspec	Enter address family ipv4 flowspec mode
PE1(config-router-af)# neighbor 6.6.6.6 activate	Activate the neighbour
PE1(config-router-af)# neighbor 10.1.1.1 activate	Activate the neighbour
PE1(config-router-af)# neighbor 20.1.1.2 activate	Activate the neighbour
PE1(config-router-af)# exit-address-family	Exit from address family
PE1(config-router)# address-family vpnv4 flowspec	Enter address family vpnv4 flowspec mode
PE1(config-router-af)# neighbor 6.6.6.6 activate	Activate the neighbour
PE1(config-router-af)# neighbor 20.1.1.2 activate	Activate the neighbour
PE1(config-router-af)# exit-address-family	Exit from address family
PE1(config-router)# address-family ipv4 vrf vrf100	Enter address family ipv4 vrf vrf name mode
PE1(config-router-af)# redistribute connected	Enter redistribute connected
PE1(config-router-af)# neighbor 66.1.1.1 remote-as 100	Enter neighbor ip add with remote-as 100
PE1(config-router-af)# neighbor 66.1.1.1 activate	Activate the neighbour
PE1(config-router-af)# exit-address-family	Exit from address family
PE1(config-router)# address-family ipv4 flowspec vrf vrf100	Enter address family ipv4 flowspec vrf vrf name mode

© 2026 IP Infusion Inc. Proprietary

PE1(config-router-af)# neighbor 66.1.1.1 activate	Activate the neighbour
PE1(config-router-af)# exit-address-family	Exit from address family
PE1(config-router)#exit	Exit from router bgp instance
PE1(config)#commit	Commits the configuration

RR2 Configuration

RR2#configure terminal	Enter configure mode.
RR2(config)#router ldp	Enter router ldp
RR2(config-router)#exit	Exit from ldp mode
RR2(config)#commit	Commits the configuration
RR2(config)#interface lo	Enter interface mode
RR2(config-if)#ip address 2.2.2.2/32 secondary	Assigned ip add as secondary
RR2(config-if)# ip router isis 1	Enter isis instance
RR2(config-if)#exit	Exit from interface mode
RR2(config)#commit	Commits the configuration
RR2(config)#interface ce43	Enter interface mode
RR2(config-if)# speed 40g	Set the speed
RR2(config-if)# ip address 11.1.1.2/24	Assign the ip add
RR2(config-if)# label-switching	Enable label switching
RR2(config-if)# ip router isis 1	Enter router isis instance
RR2(config-if)# enable-ldp ipv4	Enable ldp ipv4
RR2(config-if)#exit	Exit from interface mode
RR2(config)#commit	Commits the configuration
RR2(config)#interface xe7	Enter interface mode
RR2(config-if)# ip address 20.1.1.2/24	Assign the ip add
RR2(config-if)# label-switching	Enable label switching
RR2(config-if)# ip router isis 1	Enter router isis instance
RR2(config-if)# enable-ldp ipv4	Enable ldp ipv4
RR2(config-if)#exit	Exit from interface mode
RR2(config)#commit	Commits the configuration
RR2(config)#router isis 1	Configure isis instance
RR2(config-router)# is-type level-1-2	Configure is-type-level-1-2
RR2(config-router)# metric-style wide	Configure metric style wide
RR2(config-router)# mpls traffic-eng router-id 2.2.2.2	Configure router id
RR2(config-router)# mpls traffic-eng level-1	Configure traffic engg level-1
RR2(config-router)# mpls traffic-eng level-2	Configure traffic engg level-2

© 2026 IP Infusion Inc. Proprietary

`RR2(config-router)# dynamic-hostname`	Configure hostname
`RR2(config-router)# bfd all-interfaces`	Enable bfd all interfaces
`RR2(config-router)# net 49.0000.0000.0002.00`	Configure network id
`RR2(config-router)#exit`	Exit from isis instance mode
`RR2(config)#commit`	Commits the configuration
`RR2(config)#router bgp 65000`	Configure router bgp instance
`RR2(config-router)# no bgp inbound-route-filter`	Configure no bgp inbound route filter
`RR2(config-router)# neighbor 11.1.1.1 remote-as 65000`	Configure neighbor ip address remote as 65000
`RR2(config-router)# neighbor 20.1.1.1 remote-as 65000`	Configure neighbor ip address remote as 65000
`RR2(config-router)# address-family ipv4 unicast`	Enter address family ipv4 unicast mode
`RR2(config-router-af)# neighbor 11.1.1.1 activate`	Activate the neighbour
`RR2(config-router-af)# neighbor 11.1.1.1 route-reflector-client`	Enable the neighbor ip add as route reflector client
`RR2(config-router-af)# neighbor 20.1.1.1 activate`	Activate the neighbour
`RR2(config-router-af)# neighbor 20.1.1.1 route-reflector-client`	Enable the neighbor ip add as route reflector client
`RR2(config-router-af)# exit-address-family`	Exit from address family
`RR2(config-router)# address-family vpnv4 unicast`	Enter address family vpnv4 unicast mode
`RR2(config-router-af)# neighbor 11.1.1.1 activate`	Activate the neighbour
`RR2(config-router-af)# neighbor 11.1.1.1 route-reflector-client`	Enable the neighbor ip add as route reflector client
`RR2(config-router-af)# neighbor 20.1.1.1 activate`	Activate the neighbour
`RR2(config-router-af)# neighbor 20.1.1.1 route-reflector-client`	Enable the neighbor ip add as route reflector client
`RR2(config-router-af)# exit-address-family`	Exit from address family
`RR2(config-router)# address-family ipv4 flowspec`	Enter the address family ipv4 flowspec mode
`RR2(config-router-af)# neighbor 11.1.1.1 activate`	Activate the neighbour
`RR2(config-router-af)# neighbor 11.1.1.1 route-reflector-client`	Enable the neighbor ip add as route reflector client
`RR2(config-router-af)# neighbor 20.1.1.1 activate`	Activate the neighbour
`RR2(config-router-af)# neighbor 20.1.1.1 route-reflector-client`	Enable the neighbor ip add as route reflector client
`RR2(config-router-af)# exit-address-family`	Exit from address family
`RR2(config-router)# address-family vpnv4 flowspec`	Enter address family vpnv4 flowspec mode
`RR2(config-router-af)# neighbor 11.1.1.1 activate`	Activate the neighbour
`RR2(config-router-af)# neighbor 11.1.1.1 route-reflector-client`	Enable the neighbor ip add as route reflector client
`RR2(config-router-af)# neighbor 20.1.1.1 activate`	Activate the neighbour
`RR2(config-router-af)# neighbor 20.1.1.1 route-`	Enable the neighbor ip add as route reflector client

reflector-client	
RR2(config-router-af)# exit-address-family	Exit from address family mode
RR2(config-router)#exit	Exit from router bgp instance mode
RR2(config)#commit	

RR1 Configuration

RR1#configure terminal	Enter configure mode.
RR1(config)#router ldp	Enter router ldp
RR1(config-router)#exit	Exit from ldp mode
RR1(config)#commit	Commits the configuration
RR1(config)#interface lo	Enter the interface mode
RR1(config-if)# ip address 6.6.6.6/32 secondary	Assign the ip address secondary
RR1(config-if)# ip router isis 1	Enter router isis instance
RR1(config-if)#exit	Exit from interface mode
RR1(config)#commit	Commits the configuration
RR1(config)#interface ce4	Enter the interface mode
RR1(config-if)# ip address 50.1.1.2/24	Assign the ip add
RR1(config-if)# label-switching	Enable label switching
RR1(config-if)# ip router isis 1	Enter router isis instance
RR1(config-if)# enable-ldp ipv4	Enable ldp ipv4
RR1(config-if)#exit	Exit from interface mode
RR1(config)#commit	Commits the configuration
RR1(config)#interface ce6	Enter the interface mode
RR1(config-if)# ip address 40.1.1.2/24	Assign the ip add
RR1(config-if)# label-switching	Enable label switching
RR1(config-if)# ip router isis 1	Enter router isis instance
RR1(config-if)# enable-ldp ipv4	Enable ldp ipv4
RR1(config-if)#exit	Exit from interface mode
RR1(config)#commit	Commits the configuration
RR1(config)#router isis 1	Configure isis instance
RR1(config-router)# is-type level-1-2	Configure is-type-level-1-2
RR1(config-router)# metric-style wide	Configure metric style wide
RR1(config-router)# mpls traffic-eng router-id 6.6.6.6	Configure router id
RR1(config-router)# mpls traffic-eng level-1	Configure traffic engg level-1
RR1(config-router)# mpls traffic-eng level-2	Configure traffic engg level-2
RR1(config-router)# dynamic-hostname	Configure hostname

© 2026 IP Infusion Inc. Proprietary

`RR1(config-router)# bfd all-interfaces`	Enable bfd all interfaces
`RR1(config-router)# net 49.0000.0000.0006.00`	Configure network id
`RR1(config-router)#exit`	Exit from isis instance mode
`RR1(config)#commit`	Commits the configuration
`RR1(config)#router bgp 65000`	Configure router bgp instance
`RR1(config-router)# no bgp inbound-route-filter`	Configure no bgp inbound route filter
`RR1(config-router)# neighbor 1.1.1.1 remote-as 65000`	Configure neighbor ip address remote as 65000
`RR1(config-router)# neighbor 3.3.3.3 remote-as 65000`	Configure neighbor ip address remote as 65000
`RR1(config-router)# neighbor 1.1.1.1 update-source lo`	Configure neighbor ip address with update loopback
`RR1(config-router)# neighbor 3.3.3.3 update-source lo`	Configure neighbor ip address with update loopback
`RR1(config-router)# address-family ipv4 unicast`	Enter address family ipv4 unicast mode
`RR1(config-router-af)# neighbor 1.1.1.1 activate`	Activate the neighbour
`RR1(config-router-af)# neighbor 1.1.1.1 route-reflector-client`	Enable the neighbor ip add as route reflector client
`RR1(config-router-af)# neighbor 3.3.3.3 activate`	Activate the neighbour
`RR1(config-router-af)# neighbor 3.3.3.3 route-reflector-client`	Enable the neighbor ip add as route reflector client
`RR1(config-router-af)# exit-address-family`	Exit from address family
`RR1(config-router)# address-family vpnv4 unicast`	Enter address family vpnv4 unicast mode
`RR1(config-router-af)# neighbor 1.1.1.1 activate`	Activate the neighbour
`RR1(config-router-af)# neighbor 1.1.1.1 route-reflector-client`	Enable the neighbor ip add as route reflector client
`RR1(config-router-af)# neighbor 3.3.3.3 activate`	Activate the neighbour
`RR1(config-router-af)# neighbor 3.3.3.3 route-reflector-client`	Enable the neighbor ip add as route reflector client
`RR1(config-router-af)# exit-address-family`	Exit from address family
`RR1(config-router)# address-family ipv4 flowspec`	Enter the address family ipv4 flowspec mode
`RR1(config-router-af)# neighbor 1.1.1.1 activate`	Activate the neighbour
`RR1(config-router-af)# neighbor 1.1.1.1 route-reflector-client`	Enable the neighbor ip add as route reflector client
`RR1(config-router-af)# neighbor 3.3.3.3 activate`	Activate the neighbour
`RR1(config-router-af)# neighbor 3.3.3.3 route-reflector-client`	Enable the neighbor ip add as route reflector client
`RR1(config-router-af)# exit-address-family`	Exit from address family
`RR1(config-router)# address-family vpnv4 flowspec`	Enter address family vpnv4 flowspec mode
`RR1(config-router-af)# neighbor 1.1.1.1 activate`	Activate the neighbour
`RR1(config-router-af)# neighbor 1.1.1.1 route-reflector-client`	Enable the neighbor ip add as route reflector client

© 2026 IP Infusion Inc. Proprietary

`RR1(config-router-af)# neighbor 3.3.3.3 activate`	Activate the neighbour
`RR1(config-router-af)# neighbor 3.3.3.3 route-reflector-client`	Enable the neighbor ip add as route reflector client
`RR1(config-router-af)# exit-address-family`	Exit from address family mode
`RR1(config-router)#exit`	Exit from router bgp instance mode
`RR1(config)#commit`	

PE2 Configuration

`PE2#configure terminal`	Enter configure mode.
`PE2(config)#router ldp`	Enter vrf vrf name
`PE2(config-router)# targeted-peer ipv4 1.1.1.1`	Enter rd value
`PE2(config-router-targeted-peer)#exit-targeted-peer-mode`	Enter router targeted as both
`PE2(config-router)#exit`	Exit from vrf mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#ip vrf vrf100`	Enter vrf vrf name
`PE2(config-vrf)# rd 100:1`	Enter rd value
`PE2(config-vrf)# route-target both 100:1`	Enter router targeted as both
`PE2(config-vrf)#exit`	Exit from vrf mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#ip vrf vrf101`	Enter vrf vrf name
`PE2(config-vrf)# rd 101:1`	Enter rd value
`PE2(config-vrf)# route-target both 101:1`	Enter router targeted as both
`PE2(config-vrf)#exit`	Exit from vrf mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#class-map type traffic match-all cl1`	Configure class map with class name
`PE2(config-cmap-tr)# match destination-address ipv4 1.1.1.1/32`	Configure match with destination ip address
`PE2(config-cmap-tr)# match source-address ipv4 3.3.3.3/32`	Enable match with source ip address
`PE2(config-cmap-tr)# match port 111-222`	Enable match with port range
`PE2(config-cmap-tr)# match destination-port 300-399`	Enable match with destination-port range
`PE2(config-cmap-tr)# match source-port 230-299`	Enable match with source-port range
`PE2(config-cmap-tr)#exit`	Exit from class map traffic
`PE2(config)#commit`	Commits the configuration
`PE2(config)#policy-map type pbr pol1`	Configure policy map type pbr with policy name
`PE2(config-pmap-pbr)# class type traffic cl1`	Enable class type traffic with class map name
`PE2(config-pmap-pbr-c)# police rate 1111111 bps`	Enable policy rate

© 2026 IP Infusion Inc. Proprietary

`PE2(config-pmap-pbr-c)#exit`	Exit from polocy map traffic
`PE2(config-pmap-pbr)#exit`	Exit from policy map
`PE2(config)#commit`	Commits the configuration
`PE2(config)#flowspec`	Configure flowspec
`PE2(config-flowspec)# address-family ipv4`	Enter address-family mode with ipv4
`PE2(config-flowspec-af)# service-policy type pbr pol1`	Enable service plocy type pbr with policy name
`PE2(config-flowspec-af)#exit`	Exit from address family mode
`PE2(config-flowspec)#exit`	Exit from flowspec mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#interface lo`	Enter the loopback interface
`PE2(config-if)# ip address 3.3.3.3/32 secondary`	Assign the ip add secondary
`PE2(config-if)# ip router isis 1`	Enable router isis instance
`PE2(config-if)#exit`	Exit from interface mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#interface ce4`	Enter the interface mode
`PE2(config-if)# ip address 50.1.1.1/24`	Assign the ip add
`PE2(config-if)# label-switching`	Enable label switching
`PE2(config-if)# ip router isis 1`	Enter router isis instance
`PE2(config-if)# enable-ldp ipv4`	Enable ldp ipv4
`PE2(config-if)#exit`	Exit from interface mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#interface ce43`	Enter the interface mode
`PE2(config-if)# speed 40g`	Set the interface speed
`PE2(config-if)# ip address 11.1.1.1/24`	Assign the ip add
`PE2(config-if)# label-switching`	Enable label switching
`PE2(config-if)# ip router isis 1`	Enter router isis instance
`PE2(config-if)#enable-ldp ipv4`	Enable ldp ipv4
`PE2(config-if)#exit`	Exit from interface mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#interface ce6`	Enter the interface mode
`PE2(config-if)# ip address 30.1.1.2/24`	Assign the ip add
`PE2(config-if)#exit`	Exit from interface mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#interface ce6.10`	Enter the sub-interface mode
`PE2(config-if)# encapsulation dot1q 10`	Enable encapsulate dot1q vlan id

© 2026 IP Infusion Inc. Proprietary

`PE2(config-if)# ip vrf forwarding vrf100`	Enable brf forwarding vrf name
`PE2(config-if)# ip address 55.1.1.1/24`	Assign the ip add
`PE2(config-if)#exit`	Exit from interface mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#router isis 1`	Configure isis instance
`PE2(config-router)# is-type level-1-2`	Configure is-type-level-1-2
`PE2(config-router)# metric-style wide`	Configure metric style wide
`PE2(config-router)# mpls traffic-eng router-id 3.3.3.3`	Configure router id
`PE2(config-router)# mpls traffic-eng level-1`	Configure traffic engg level-1
`PE2(config-router)# mpls traffic-eng level-2`	Configure traffic engg level-2
`PE2(config-router)# dynamic-hostname`	Configure hostname
`PE2(config-router)# bfd all-interfaces`	Enable bfd all interfaces
`PE2(config-router)# net 49.0000.0000.0003.00`	Configure network id
`PE2(config-router)# passive-interface lo`	Configure passive interface as loopback
`PE2(config-router)#exit`	Exit from isis instance mode
`PE2(config)#commit`	Commits the configuration
`PE2(config)#router bgp 65000`	Configure bgp instance
`PE2(config-router)# neighbor 6.6.6.6 remote-as 65000`	Enable neighbour address with remote-as 65000
`PE2(config-router)# neighbor 11.1.1.2 remote-as 65000`	Enable neighbour address with remote-as 65000
`PE2(config-router)# neighbor 30.1.1.1 remote-as 200`	Enable neighbour address with remote-as 300
`PE2(config-router)# neighbor 6.6.6.6 update-source lo`	Enable neighbour address with update loopback
`PE2(config-router)# address-family ipv4 unicast`	Enter address family ipv4 mode
`PE2(config-router-af)# neighbor 6.6.6.6 activate`	Activate the neighbour
`PE2(config-router-af)# neighbor 6.6.6.6 next-hop-self`	Enter neighbor ip add with next hop self
`PE2(config-router-af)# neighbor 11.1.1.2 activate`	Activate the neighbour
`PE2(config-router-af)# neighbor 11.1.1.2 next-hop-self`	Activate the neighbour
`PE2(config-router-af)# neighbor 30.1.1.1 activate`	Enter neighbor ip add with next hop self
`PE2(config-router-af)# exit-address-family`	Exit from address family
`PE2(config-router)# address-family vpnv4 unicast`	Enter address family vpnv4 unicast mode
`PE2(config-router-af)# neighbor 6.6.6.6 activate`	Activate the neighbour
`PE2(config-router-af)# neighbor 6.6.6.6 next-hop-self`	Enter neighbor ip add with next hop self
`PE2(config-router-af)# neighbor 11.1.1.2 activate`	Activate the neighbour
`PE2(config-router-af)# neighbor 11.1.1.2 next-hop-`	Enter neighbor ip add with next hop self

© 2026 IP Infusion Inc. Proprietary

self	
PE2(config-router-af)# exit-address-family	Exit from address family
PE2(config-router)# address-family ipv4 flowspec	Enter address family ipv4 flowspec mode
PE2(config-router-af)# neighbor 6.6.6.6 activate	Activate the neighbour
PE2(config-router-af)# neighbor 11.1.1.2 activate	Activate the neighbour
PE2(config-router-af)# neighbor 30.1.1.1 activate	Activate the neighbour
PE2(config-router-af)# exit-address-family	Exit from address family
PE2(config-router)# address-family vpnv4 flowspec	Enter address family vpnv4 flowspec mode
PE2(config-router-af)# neighbor 6.6.6.6 activate	Activate the neighbour
PE2(config-router-af)# neighbor 11.1.1.2 activate	Activate the neighbour
PE2(config-router-af)# exit-address-family	Exit from address family
PE2(config-router)# address-family ipv4 vrf vrf100	Enter address family ipv4 vrf vrf name mode
PE2(config-router-af)# redistribute connected	Enter redistribute connected
PE2(config-router-af)# neighbor 55.1.1.2 remote-as 200	Enter neighbor ip add with remote-as 200
PE2(config-router-af)# neighbor 55.1.1.2 activate	Activate the neighbour
PE2(config-router-af)# exit-address-family	Exit from address family
PE2(config-router)# address-family ipv4 flowspec vrf vrf100	Enter address family ipv4 flowspec vrf vrf name mode
PE2(config-router-af)# neighbor 55.1.1.2 activate	Activate the neighbour
PE2(config-router-af)# exit-address-family	Exit from address family
PE2(config-router)#exit	Exit from router bgp instance
PE2(config)#commit	Commits the configuration

CE2 Configuration

CE2#configure terminal	Enter configure mode.
CE2(config)#class-map type traffic match-all clas1	Configure class map with class name
CE2(config-cmap-tr)# match destination-address ipv4 4.4.4.4/32	Configure match with destination ip address
CE2(config-cmap-tr)# match source-address ipv4 5.5.5.5/32	Enable match with source ip address
CE2(config-cmap-tr)# match protocol 121	Enable match with protocol
CE2(config-cmap-tr)# match port 777-888	Enable match with port
CE2(config-cmap-tr)# match destination-port 1300-1400	Enable match with destination-port
CE2(config-cmap-tr)# match source-port 1001-1200	Enable match with source-port
CE2(config-cmap-tr)#exit	Exit from class map traffic mode
CE2(config)#commit	Commits the configuration
CE2(config)#class-map type traffic match-all clas2	Configure class map with class name

© 2026 IP Infusion Inc. Proprietary

`CE2(config-cmap-tr)# match destination-address ipv4 30.1.1.0/24`	Configure match with destination ip address
`CE2(config-cmap-tr)# match source-address ipv4 30.1.1.0/24`	Enable match with source ip address
`CE2(config-cmap-tr)# match protocol 11`	Enable match with protocol
`CE2(config-cmap-tr)# match port 100-200`	Enable match with port
`CE2(config-cmap-tr)# match destination-port 300-400`	Enable match with destination-port
`CE2(config-cmap-tr)# match source-port 500-600`	Enable match with source-port
`CE2(config-cmap-tr)# match tcp-flag 21 bit-mask 3f`	Enable match with tcp flag bit
`CE2(config-cmap-tr)# match dscp 21`	Enable match with dscp
`CE2(config-cmap-tr)#exit`	Exit from class map traffic mode
`CE2(config)#commit`	Commits the configuration
`CE2(config)#policy-map type pbr pol1`	Configure policy map type pbr with policy name
`CE2(config-pmap-pbr)#class type traffic clas1`	Enable class type traffic with class map name
`CE2(config-pmap-pbr-c)#police rate 100000 bps`	Enable policy rate
`CE2(config-pmap-pbr-c)#exit`	Exit from polocy map traffic
`CE2(config-pmap-pbr)#exit`	Exit from policy map
`CE2(config)#commit`	Commits the configuration
`CE2(config)#policy-map type pbr pol2`	Configure policy map type pbr with policy name
`CE2(config-pmap-pbr)# class type traffic clas2`	Enable class type traffic with class map name
`CE2(config-pmap-pbr-c)# police rate 500000 bps`	Enable policy rate
`CE2(config-pmap-pbr-c)# set dscp 1e`	Ser the dscp valu
`CE2(config-pmap-pbr-c)#exit`	Exit from polocy map traffic
`CE2(config-pmap-pbr)#exit`	Exit from policy map
`CE2(config)#commit`	Commits the configuration
`CE2(config)#flowspec`	Configure flowspec
`CE2(config-flowspec)# address-family ipv4`	Enter address-family mode with ipv4
`CE2(config-flowspec-af)# service-policy type pbr pol1`	Enable service plocy type pbr with policy name
`CE2(config-flowspec-af)# service-policy type pbr pol2`	Enable service plocy type pbr with policy name
`CE2(config-flowspec-af)#exit`	Exit from address family mode
`CE2(config-flowspec)#exit`	Exit from flowspec mode
`CE2(config)#commit`	Commits the configuration
`CE2(config)#interface lo`	Enter interface mode
`CE2(config-if)# ip address 5.5.5.5/32 secondary`	Assign the ip address secondary
`CE2(config-if)# ip router isis 1`	Enter router isis instance
`CE2(config-if)#exit`	Exit from interface mode

`CE2(config)#commit`	Commits the configuration
`CE2(config)#interface ce49`	Enter interface mode
`CE2(config-if)# ip address 30.1.1.1/24`	Assign the ip address
`CE2(config-if)#exit`	Exit from interface mode
`CE2(config)#commit`	Commits the configuration
`CE2-7025(config)#interface ce49.10`	Enter sub-interface mode
`CE2-7025(config-if)# encapsulation dot1q 10`	Enable encapsulate dot1q vlan id
`CE2-7025(config-if)# ip address 55.1.1.2/24`	Assign the ip address
`CE2-7025(config-if)#exit`	Exit from interface mode
`CE2-7025(config)#commit`	Commits the configuration
`CE2(config)#router bgp 200`	Configure bgp instance
`CE2(config-router)# neighbor 30.1.1.2 remote-as 65000`	Enable neighbour address with remote-as 65000
`CE2(config-router)# neighbor 55.1.1.1 remote-as 65000`	Enable neighbour address with remote-as 65000
`CE2(config-router)# address-family ipv4 unicast`	Enter address family ipv4 unicast mode
`CE2(config-router-af)# redistribute connected`	Enter redistribute connected
`CE2(config-router-af)# neighbor 30.1.1.2 activate`	Activate the neighbour
`CE2(config-router-af)# neighbor 55.1.1.1 activate`	Activate the neighbour
`CE2(config-router-af)# exit-address-family`	Exit from address family mode
`CE2(config-router)# address-family ipv4 flowspec`	Enter address family ipv4 flowspec mode
`CE2(config-router-af)# neighbor 30.1.1.2 activate`	Activate the neighbour
`CE2(config-router-af)# neighbor 55.1.1.1 activate`	Exit from address family
`CE2(config-router-af)# exit-address-family`	Exit from router bgp mode
`CE2(config-router)#exit`	Exit from router bgp instance
`CE2(config)#commit`	Commits the configuration

BGP Flowspec Configuration With RR Validation

CE1

```
CE1#sh ip bgp
BGP table version is 3, local router ID is 4.4.4.4
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
*>   4.4.4.4/32       0.0.0.0               0        100       32768  ?
*>   5.5.5.5/32       10.1.1.2              0        100       0   65000 200 ?
*                     66.1.1.2              0        100       0   65000 200 ?
*>   10.1.1.0/24      0.0.0.0               0        100       32768  ?
*>   30.1.1.0/24      10.1.1.2              0        100       0   65000 200 ?
*                     66.1.1.2              0        100       0   65000 200 ?
*>   55.1.1.0/24      66.1.1.2              0        100       0   65000 ?
*                     10.1.1.2              0        100       0   65000 200 ?
*>   66.1.1.0/24      0.0.0.0               0        100       32768  ?
```

```
*                     66.1.1.2            0        100       0   65000 ?
*>   100.1.1.0/24     0.0.0.0             0        100       32768  ?
*>   200.1.1.0        10.1.1.2            0        100       0   65000 200 ?
*                     66.1.1.2            0        100       0   65000 200 ?

Total number of prefixes 8
CE1#

CE1#sh ip bgp neighbors
BGP neighbor is 10.1.1.2, remote AS 65000, local AS 100, external link
  BGP version 4, local router ID 4.4.4.4, remote router ID 1.1.1.1
  BGP state = Established, up for 03:06:16
  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
    Address family IPv4 Flow Specification: advertised and received
  Received 447 messages, 0 notifications, 0 in queue
  Sent 444 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 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  5 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 2, neighbor version 2
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 10.1.1.1, Local port: 179
Foreign host: 10.1.1.2, Foreign port: 37390
Nexthop: 10.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 03:06:16, due to OPEN Message Error (Notification sent)
Notification Error Message: (OPEN Message Error/Bad BGP Identifier. Reason: )

BGP neighbor is 66.1.1.2, remote AS 65000, local AS 100, external link
  BGP version 4, local router ID 4.4.4.4, remote router ID 66.1.1.2
  BGP state = Established, up for 03:06:15
  Last read 00:00:07, 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 IPv4 Flow Specification: advertised and received
  Received 446 messages, 0 notifications, 0 in queue
  Sent 446 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 3, neighbor version 3
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  5 accepted prefixes
  7 announced prefixes
```

```
For address family: IPv4 Flow Specification  BGP table version 2, neighbor version 2
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 66.1.1.1, Local port: 53648
Foreign host: 66.1.1.2, Foreign port: 179
Nexthop: 66.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 03:06:20, due to OPEN Message Error (Notification sent)
Notification Error Message: (OPEN Message Error/Bad BGP Identifier. Reason: )

CE1#
CE1#
CE1#sh ip bgp ipv4 flowspec summary
BGP router identifier 4.4.4.4, local AS number 100
BGP table version is 2
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
10.1.1.2                 4 65000  447        445       2      0      0  03:06:20               4
66.1.1.2                 4 65000  446        448       2      0      0  03:06:19               2

Total number of neighbors 2

Total number of Established sessions 2
CE1#
CE1#sh ip bgp ipv4 flowspec
BGP table version is 2, local router ID is 4.4.4.4
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
*>   flowspec/216     0.0.0.0               0        100       32768  i
*    flowspec/336     10.1.1.2              0        100       0   65000 200 i
*                     66.1.1.2              0        100       0   65000 200 i
*    flowspec/264     10.1.1.2              0        100       0   65000 i
*    flowspec/392     10.1.1.2              0        100       0   65000 i
*>   flowspec/176     0.0.0.0               0        100       32768  i
*    flowspec/288     10.1.1.2              0        100       0   65000 200 i
*                     66.1.1.2              0        100       0   65000 200 i

Total number of prefixes 6
CE1#
CE1#sh ip bgp ipv4 flowspec detail
BGP table version is 2, local router ID is 4.4.4.4
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
*>   flowspec/216     0.0.0.0               0        100       32768  i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:44
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144
```

```
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*    flowspec/336     10.1.1.2             0        100       0   65000 200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:11
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:100
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:200
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*                    66.1.1.2             0        100       0   65000 200 i
  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*    flowspec/264     10.1.1.2             0        100       0   65000 i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:111
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:230
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:299

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*    flowspec/392     10.1.1.2             0        100       0   65000 i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 1.1.1.1/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:78
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:101
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:201
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:400
```

© 2026 IP Infusion Inc. Proprietary

```
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:499
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:330
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    Packet length:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x10

  FlowSpec Actions:
    traffic_rate: 1000011 bps
    traffic marking DSCP: e

*>   flowspec/176     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 4.4.4.4/32
    source ip address: 4.4.4.4/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:12
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:333
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:444

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*    flowspec/288     10.1.1.2             0        100       0   65000 200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:121
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:777
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:888
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1001
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1200

  FlowSpec Actions:
    traffic_rate: 100000 bps

*                    66.1.1.2             0        100       0   65000 200 i
  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 6
CE1#
CE1#
CE1#sh ip bgp ipv4 flowspec vrf all
BGP table version is 2, local router ID is 4.4.4.4
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
*>   flowspec/216     0.0.0.0              0        100       32768  i
*    flowspec/336     10.1.1.2             0        100       0   65000 200 i
*                     66.1.1.2             0        100       0   65000 200 i
```

```
*    flowspec/264     10.1.1.2             0        100       0   65000 i
*    flowspec/392     10.1.1.2             0        100       0   65000 i
*>   flowspec/176     0.0.0.0              0        100       32768  i
*    flowspec/288     10.1.1.2             0        100       0   65000 200 i
*                     66.1.1.2             0        100       0   65000 200 i

Total number of prefixes 6
CE1#sh ip bgp ipv4 flowspec vrf all detail
BGP table version is 2, local router ID is 4.4.4.4
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
*>   flowspec/216     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:44
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*    flowspec/336     10.1.1.2             0        100       0   65000 200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:11
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:100
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:200
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*                    66.1.1.2             0        100       0   65000 200 i
  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*    flowspec/264     10.1.1.2             0        100       0   65000 i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP port:
```

```
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:111
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:230
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:299

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*    flowspec/392     10.1.1.2             0        100       0   65000 i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 1.1.1.1/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:78
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:101
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:201
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:400
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:499
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:330
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    Packet length:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x10

  FlowSpec Actions:
    traffic_rate: 1000011 bps
    traffic marking DSCP: e

*>   flowspec/176     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 4.4.4.4/32
    source ip address: 4.4.4.4/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:12
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:333
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:444

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*    flowspec/288     10.1.1.2             0        100       0   65000 200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:121
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:777
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:888
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400
    TCP/UDP source port:
```

```
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1001
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1200

  FlowSpec Actions:
    traffic_rate: 100000 bps

*                     66.1.1.2             0        100       0   65000 200 i
  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 6
CE1#
```

PE1

```
PE1#sh ip bgp
BGP table version is 5, local router ID is 1.1.1.1
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
*>   4.4.4.4/32       10.1.1.1              0        100       0   100 ?
*>i  5.5.5.5/32       3.3.3.3               0        100       0   200 ?
* i                   11.1.1.1              0        100       0   200 ?
*    9.9.9.9/32       0.0.0.0               0        100       32768  i
*>   10.1.1.0/24      10.1.1.1              0        100       0   100 ?
*>i  30.1.1.0/24      3.3.3.3               0        100       0   200 ?
* i                   11.1.1.1              0        100       0   200 ?
*>i  55.1.1.0/24      3.3.3.3               0        100       0   200 ?
* i                   11.1.1.1              0        100       0   200 ?
*>   66.1.1.0/24      10.1.1.1              0        100       0   100 ?
*>   100.1.1.0/24     10.1.1.1              0        100       0   100 ?
*>i  200.1.1.0        3.3.3.3               0        100       0   200 ?
* i                   11.1.1.1              0        100       0   200 ?

Total number of prefixes 9
PE1#
PE1#sh ip bgp neighbors
BGP neighbor is 6.6.6.6, remote AS 65000, local AS 65000, internal link
  BGP version 4, local router ID 1.1.1.1, remote router ID 6.6.6.6
  BGP state = Established, up for 03:06:53
  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
    Address family VPNv4 Unicast: advertised and received
    Address family IPv4 Flow Specification: advertised and received
    Address family VPNv4 Flow Specification: advertised and received
  Received 448 messages, 0 notifications, 0 in queue
  Sent 445 messages, 1 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 5, neighbor version 5
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: VPNv4 Unicast  BGP table version 3, neighbor version 3
```

```
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 7, neighbor version 7
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  3 accepted prefixes
  3 announced prefixes


For address family: VPNv4 Flow Specification  BGP table version 2, neighbor version 2
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 1.1.1.1, Local port: 41307
Foreign host: 6.6.6.6, Foreign port: 179
Nexthop: 1.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 03:06:58, due to OPEN Message Error (Notification sent)
Notification Error Message: (OPEN Message Error/Bad BGP Identifier. Reason: )

BGP neighbor is 10.1.1.1, remote AS 100, local AS 65000, external link
  BGP version 4, local router ID 1.1.1.1, remote router ID 4.4.4.4
  BGP state = Established, up for 03:07:28
  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
    Address family IPv4 Flow Specification: advertised and received
  Received 444 messages, 1 notifications, 0 in queue
  Sent 450 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 5, neighbor version 5
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 7, neighbor version 7
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  4 announced prefixes

 Connections established 1; dropped 0
Local host: 10.1.1.2, Local port: 37390
Foreign host: 10.1.1.1, Foreign port: 179
Nexthop: 10.1.1.2
Nexthop global: ::
Nexthop local: ::
```

```
BGP connection: non shared network
Last Reset: 03:07:33, due to BGP Notification received
Notification Error Message: (OPEN Message Error/Bad BGP Identifier. Reason: )

BGP neighbor is 20.1.1.2, remote AS 65000, local AS 65000, internal link
  BGP version 4, local router ID 1.1.1.1, remote router ID 2.2.2.2
  BGP state = Established, up for 03:06:26
  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
    Address family VPNv4 Unicast: advertised and received
    Address family IPv4 Flow Specification: advertised and received
    Address family VPNv4 Flow Specification: advertised and received
  Received 446 messages, 0 notifications, 0 in queue
  Sent 450 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 5, neighbor version 5
  Index 3, Offset 0, Mask 0x8
  AIGP is enabled
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: VPNv4 Unicast  BGP table version 3, neighbor version 3
  Index 2, Offset 0, Mask 0x4
  AIGP is enabled
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 7, neighbor version 7
  Index 3, Offset 0, Mask 0x8
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  3 accepted prefixes
  3 announced prefixes


For address family: VPNv4 Flow Specification  BGP table version 2, neighbor version 2
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 20.1.1.1, Local port: 179
Foreign host: 20.1.1.2, Foreign port: 43160
Nexthop: 20.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 03:06:26, due to Administratively Reset (Cease Notification sent)
Notification Error Message: (Cease/Administratively Reset.)

BGP neighbor is 66.1.1.1, vrf vrf100, remote AS 100, local AS 65000, external link
  BGP version 4, local router ID 66.1.1.2, remote router ID 4.4.4.4
  BGP state = Established, up for 03:07:28
  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
    Address family IPv4 Flow Specification: advertised and received
  Received 449 messages, 1 notifications, 0 in queue
  Sent 451 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)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  5 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 66.1.1.2, Local port: 179
Foreign host: 66.1.1.1, Foreign port: 53648
Nexthop: 66.1.1.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 03:07:33, due to BGP Notification received
Notification Error Message: (OPEN Message Error/Bad BGP Identifier. Reason: )

PE1#
PE1#
PE1#sh ip bgp ipv4 flowspec summary
BGP router identifier 1.1.1.1, local AS number 65000
BGP table version is 7
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
6.6.6.6                  4 65000  448        446       7      0      0  03:07:01               3
10.1.1.1                 4   100  446        451       7      0      0  03:07:36               2
20.1.1.2                 4 65000  446        452       7      0      0  03:06:34               3

Total number of neighbors 3

Total number of Established sessions 3
BGP router identifier 66.1.1.2, local AS number 65000
BGP VRF vrf100 Route Distinguisher: 100:1
BGP table version is 1
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
66.1.1.1                 4   100  451        451       1      0      0  03:07:36               2

Total number of neighbors 1

Total number of Established sessions 1
PE1#
PE1#
PE1#sh ip bgp ipv4 flowspec
BGP table version is 7, local router ID is 1.1.1.1
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
*>   flowspec/216     10.1.1.1              0        100       0   100 i
*>i  flowspec/336     6.6.6.6               0        100       0   200 i
* i                   20.1.1.2              0        100       0   200 i
*>i  flowspec/264     6.6.6.6               0        100       0    i
* i                   20.1.1.2              0        100       0    i
*>   flowspec/392     0.0.0.0               0        100       32768  i
*>   flowspec/176     10.1.1.1              0        100       0   100 i
*>i  flowspec/288     6.6.6.6               0        100       0   200 i
* i                   20.1.1.2              0        100       0   200 i

Total number of prefixes 6
PE1#
PE1#sh ip bgp ipv4 flowspec detail
BGP table version is 7, local router ID is 1.1.1.1
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
*>   flowspec/216     10.1.1.1              0        100       0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:44
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>i  flowspec/336     6.6.6.6               0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:11
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:100
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:200
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

* i                   20.1.1.2              0        100       0   200 i
  FlowSpec Actions:
    traffic_rate: 500000 bps
```

```
    traffic marking DSCP: 1e

*>i  flowspec/264     6.6.6.6              0        100       0    i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:111
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:230
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:299

  FlowSpec Actions:
    traffic_rate: 1111111 bps

* i                   20.1.1.2             0        100       0    i
  FlowSpec Actions:
    traffic_rate: 1111111 bps

*>   flowspec/392     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 1.1.1.1/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:78
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:101
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:201
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:400
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:499
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:330
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    Packet length:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x10

  FlowSpec Actions:
    traffic_rate: 1000011 bps
    traffic marking DSCP: e

*>   flowspec/176     10.1.1.1             0        100       0   100 i

  FlowSpec:
    destination ip address: 4.4.4.4/32
    source ip address: 4.4.4.4/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:12
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:333
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:444

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>i  flowspec/288     6.6.6.6              0        100       0   200 i
```

```
  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:121
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:777
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:888
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1001
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1200

  FlowSpec Actions:
    traffic_rate: 100000 bps

* i                  20.1.1.2             0        100       0   200 i
  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 6
PE1#
PE1#
PE1#sh ip bgp ipv4 flowspec vrf all
BGP table version is 1, local router ID is 66.1.1.2
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

BGP Route Table for VRF vrf100
*>   flowspec/216     66.1.1.1             0        100       0   100 i
*>i  flowspec/336     6.6.6.6              0        100       0   200 i
*>   flowspec/176     66.1.1.1             0        100       0   100 i
*>i  flowspec/288     6.6.6.6              0        100       0   200 i

Total number of prefixes 4
*>   flowspec/216     10.1.1.1             0        100       0   100 i
*>i  flowspec/336     6.6.6.6              0        100       0   200 i
* i                   20.1.1.2             0        100       0   200 i
*>i  flowspec/264     6.6.6.6              0        100       0    i
* i                   20.1.1.2             0        100       0    i
*>   flowspec/392     0.0.0.0              0        100       32768  i
*>   flowspec/176     10.1.1.1             0        100       0   100 i
*>i  flowspec/288     6.6.6.6              0        100       0   200 i
* i                   20.1.1.2             0        100       0   200 i

Total number of prefixes 6
PE1#
PE1#
PE1#sh ip bgp ipv4 flowspec vrf all detail
BGP table version is 1, local router ID is 66.1.1.2
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

BGP Route Table for VRF vrf100
*>   flowspec/216     66.1.1.1             0        100       0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
```

```
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:44
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>i  flowspec/336     6.6.6.6              0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:11
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:100
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:200
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*>   flowspec/176     66.1.1.1             0        100       0   100 i

  FlowSpec:
    destination ip address: 4.4.4.4/32
    source ip address: 4.4.4.4/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:12
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:333
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:444

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>i  flowspec/288     6.6.6.6              0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:121
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:777
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:888
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1001
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1200
```

© 2026 IP Infusion Inc. Proprietary

```
  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 4
*>   flowspec/216     10.1.1.1             0        100       0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:44
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>i  flowspec/336     6.6.6.6              0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:11
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:100
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:200
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

* i                   20.1.1.2             0        100       0   200 i
  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*>i  flowspec/264     6.6.6.6              0        100       0    i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:111
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:230
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:299
```

© 2026 IP Infusion Inc. Proprietary

```
  FlowSpec Actions:
    traffic_rate: 1111111 bps

* i                   20.1.1.2             0        100       0    i
  FlowSpec Actions:
    traffic_rate: 1111111 bps

*>   flowspec/392     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 1.1.1.1/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:78
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:101
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:201
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:400
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:499
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:330
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    Packet length:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x10

  FlowSpec Actions:
    traffic_rate: 1000011 bps
    traffic marking DSCP: e

*>   flowspec/176     10.1.1.1             0        100       0   100 i

  FlowSpec:
    destination ip address: 4.4.4.4/32
    source ip address: 4.4.4.4/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:12
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:333
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:444

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>i  flowspec/288     6.6.6.6              0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:121
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:777
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:888
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1001
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1200

  FlowSpec Actions:
```

```
    traffic_rate: 100000 bps

* i                   20.1.1.2             0        100       0   200 i
  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 6
PE1#
PE1#
PE1#sh ip bgp vpnv4 flowspec summary
BGP router identifier 1.1.1.1, local AS number 65000
BGP table version is 2
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
6.6.6.6                  4 65000  450        448       2      0      0  03:07:39               2
20.1.1.2                 4 65000  448        453       2      0      0  03:07:12               2

Total number of neighbors 2

Total number of Established sessions 2
PE1#
PE1#sh ip bgp vpnv4 flowspec vrf vrf100
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
Route Distinguisher: 100:1 (Default for VRF vrf100)
*>   flowspec/216     66.1.1.1             0        100       0   100 i
*>i  flowspec/336     6.6.6.6              0        100       0   200 i
*>   flowspec/176     66.1.1.1             0        100       0   100 i
*>i  flowspec/288     6.6.6.6              0        100       0   200 i
 Announced routes count = 2
 Accepted routes count = 2
PE1#
PE1#sh ip bgp vpnv4 flowspec vrf vrf100 detail
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
Route Distinguisher: 100:1 (Default for VRF vrf100)
*>   flowspec/216     66.1.1.1             0        100       0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:44
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>i  flowspec/336     6.6.6.6              0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
```

```
    source ip address: 30.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:11
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:100
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:200
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*>   flowspec/176     66.1.1.1             0        100       0   100 i

  FlowSpec:
    destination ip address: 4.4.4.4/32
    source ip address: 4.4.4.4/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:12
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:333
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:444

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>i  flowspec/288     6.6.6.6              0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:121
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:777
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:888
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1001
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1200

  FlowSpec Actions:
    traffic_rate: 100000 bps

 Announced routes count = 2
 Accepted routes count = 2
PE1#
RR1
RR1#sh 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
PE2            ce4         5c07.5819.4878      Up     27        L1   IS-IS
```

```
                                          Up     27        L2   IS-IS
PE1            ce6         b86a.9725.a7ec Up     9         L1   IS-IS
                                          Up     9         L2   IS-IS
```

RR1

```
RR1#sh ip bgp
BGP table version is 3, local router ID is 6.6.6.6
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
*>i  4.4.4.4/32       1.1.1.1               0        100       0   100 ?
*>i  5.5.5.5/32       3.3.3.3               0        100       0   200 ?
*>i  10.1.1.0/24      1.1.1.1               0        100       0   100 ?
*>i  30.1.1.0/24      3.3.3.3               0        100       0   200 ?
*>i  55.1.1.0/24      3.3.3.3               0        100       0   200 ?
*>i  66.1.1.0/24      1.1.1.1               0        100       0   100 ?
*>i  100.1.1.0/24     1.1.1.1               0        100       0   100 ?
*>i  200.1.1.0        3.3.3.3               0        100       0   200 ?

Total number of prefixes 8
RR1#
RR1#
RR1#sh ip bgp neighbors
BGP neighbor is 1.1.1.1, remote AS 65000, local AS 65000, internal link
  BGP version 4, local router ID 6.6.6.6, remote router ID 1.1.1.1
  BGP state = Established, up for 03:08:38
  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
    Address family VPNv4 Unicast: advertised and received
    Address family IPv4 Flow Specification: advertised and received
    Address family VPNv4 Flow Specification: advertised and received
  Received 449 messages, 1 notifications, 0 in queue
  Sent 454 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 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: VPNv4 Unicast  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 5, neighbor version 5
  Index 1, Offset 0, Mask 0x2
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
```

```
  3 accepted prefixes
  3 announced prefixes


For address family: VPNv4 Flow Specification  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 6.6.6.6, Local port: 179
Foreign host: 1.1.1.1, Foreign port: 41307
Nexthop: 6.6.6.6
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 03:08:43, due to BGP Notification received
Notification Error Message: (OPEN Message Error/Bad BGP Identifier. Reason: )

BGP neighbor is 3.3.3.3, remote AS 65000, local AS 65000, internal link
  BGP version 4, local router ID 6.6.6.6, remote router ID 3.3.3.3
  BGP state = Established, up for 03:06:51
  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
    Address family VPNv4 Unicast: advertised and received
    Address family IPv4 Flow Specification: advertised and received
    Address family VPNv4 Flow Specification: advertised and received
  Received 446 messages, 0 notifications, 0 in queue
  Sent 446 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 3, neighbor version 3
  Index 2, Offset 0, Mask 0x4
  AIGP is enabled
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: VPNv4 Unicast  BGP table version 3, neighbor version 3
  Index 2, Offset 0, Mask 0x4
  AIGP is enabled
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 5, neighbor version 5
  Index 2, Offset 0, Mask 0x4
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  3 accepted prefixes
  3 announced prefixes


For address family: VPNv4 Flow Specification  BGP table version 3, neighbor version 3
```

```
  Index 2, Offset 0, Mask 0x4
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 6.6.6.6, Local port: 35405
Foreign host: 3.3.3.3, Foreign port: 179
Nexthop: 6.6.6.6
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

RR1#
RR1#
RR1#sh ip bgp ipv4 flowspec summary
BGP router identifier 6.6.6.6, local AS number 65000
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
1.1.1.1                  4 65000  450        454       5      0      0  03:08:46              3
3.3.3.3                  4 65000  447        447       5      0      0  03:06:59              3

Total number of neighbors 2

Total number of Established sessions 2
RR1#
RR1#
RR1#sh ip bgp ipv4 flowspec
BGP table version is 5, local router ID is 6.6.6.6
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
*>i  flowspec/216     1.1.1.1               0        100       0   100 i
*>i  flowspec/336     3.3.3.3               0        100       0   200 i
*>i  flowspec/264     3.3.3.3               0        100       0    i
*>i  flowspec/392     1.1.1.1               0        100       0    i
*>i  flowspec/176     1.1.1.1               0        100       0   100 i
*>i  flowspec/288     3.3.3.3               0        100       0   200 i

Total number of prefixes 6
RR1#
RR1#
RR1#sh ip bgp ipv4 flowspec detail
BGP table version is 5, local router ID is 6.6.6.6
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
*>i  flowspec/216     1.1.1.1               0        100       0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:44
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
```

```
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>i  flowspec/336     3.3.3.3              0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:11
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:100
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:200
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*>i  flowspec/264     3.3.3.3              0        100       0    i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:111
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:230
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:299

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*>i  flowspec/392     1.1.1.1              0        100       0    i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 1.1.1.1/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:78
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:101
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:201
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:400
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:499
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:330
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    Packet length:
```

```
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x10

  FlowSpec Actions:
    traffic_rate: 1000011 bps
    traffic marking DSCP: e

*>i  flowspec/176     1.1.1.1              0        100       0   100 i

  FlowSpec:
    destination ip address: 4.4.4.4/32
    source ip address: 4.4.4.4/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:12
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:333
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:444

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>i  flowspec/288     3.3.3.3              0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:121
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:777
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:888
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1001
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1200

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 6
RR1#
RR1#
RR1#sh ip bgp ipv4 flowspec vrf all
BGP table version is 5, local router ID is 6.6.6.6
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
*>i  flowspec/216     1.1.1.1               0        100       0   100 i
*>i  flowspec/336     3.3.3.3               0        100       0   200 i
*>i  flowspec/264     3.3.3.3               0        100       0    i
*>i  flowspec/392     1.1.1.1               0        100       0    i
*>i  flowspec/176     1.1.1.1               0        100       0   100 i
*>i  flowspec/288     3.3.3.3               0        100       0   200 i

Total number of prefixes 6
RR1#
RR1#sh ip bgp ipv4 flowspec vrf all detail
BGP table version is 5, local router ID is 6.6.6.6
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
*>i  flowspec/216     1.1.1.1               0        100       0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:44
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>i  flowspec/336     3.3.3.3               0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:11
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:100
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:200
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*>i  flowspec/264     3.3.3.3               0        100       0    i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:111
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:230
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:299

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*>i  flowspec/392     1.1.1.1               0        100       0    i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 1.1.1.1/32
```

```
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:78
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:101
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:201
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:400
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:499
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:330
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    Packet length:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x10

  FlowSpec Actions:
    traffic_rate: 1000011 bps
    traffic marking DSCP: e

*>i  flowspec/176     1.1.1.1              0        100       0   100 i

  FlowSpec:
    destination ip address: 4.4.4.4/32
    source ip address: 4.4.4.4/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:12
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:333
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:444

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>i  flowspec/288     3.3.3.3              0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:121
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:777
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:888
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1001
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1200

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 6
RR1#
RR1#
RR1#sh ip bgp vpnv4 all
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
```

```
Route Distinguisher: 100:1
*>i  4.4.4.4/32       1.1.1.1             0        100       0   100 ?
*>i  5.5.5.5/32       3.3.3.3             0        100       0   200 ?
*>i  10.1.1.0/24      1.1.1.1             0        100       0   100 ?
*>i  30.1.1.0/24      3.3.3.3             0        100       0   200 ?
*>i  55.1.1.0/24      3.3.3.3             0        100       0    ?
*>i  66.1.1.0/24      1.1.1.1             0        100       0    ?
*>i  100.1.1.0/24     1.1.1.1             0        100       0   100 ?
*>i  200.1.1.0        3.3.3.3             0        100       0   200 ?
 Announced routes count = 0
 Accepted routes count = 8
RR1#
RR1#
RR1#sh ip bgp vpnv4 flowspec summary
BGP router identifier 6.6.6.6, local AS number 65000
BGP table version is 3
3 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 65000  452        456       3      0      0  03:09:22              2
3.3.3.3                  4 65000  448        448       3      0      0  03:07:35              2

Total number of neighbors 2

Total number of Established sessions 2
RR1#
```

RR2

```
RR2#sh 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
PE1            xe7         b86a.9725.a7c2      Up     6         L1   IS-IS
                                               Up     6         L2   IS-IS
PE2            ce43        5c07.5819.486b      Up     7         L1   IS-IS
                                               Up     7         L2   IS-IS
RR2#
RR2#sh ip bgp
BGP table version is 3, local router ID is 2.2.2.2
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
*>i  4.4.4.4/32       20.1.1.1              0        100       0   100 ?
*>i  5.5.5.5/32       11.1.1.1              0        100       0   200 ?
*>i  10.1.1.0/24      20.1.1.1              0        100       0   100 ?
*>i  30.1.1.0/24      11.1.1.1              0        100       0   200 ?
*>i  55.1.1.0/24      11.1.1.1              0        100       0   200 ?
*>i  66.1.1.0/24      20.1.1.1              0        100       0   100 ?
*>i  100.1.1.0/24     20.1.1.1              0        100       0   100 ?
*>i  200.1.1.0        11.1.1.1              0        100       0   200 ?

Total number of prefixes 8
RR2#
RR2#
RR2#sh ip bgp neighbors
BGP neighbor is 11.1.1.1, remote AS 65000, local AS 65000, internal link
  BGP version 4, local router ID 2.2.2.2, remote router ID 3.3.3.3
  BGP state = Established, up for 03:09:46
  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
    Address family VPNv4 Unicast: advertised and received
    Address family IPv4 Flow Specification: advertised and received
    Address family VPNv4 Flow Specification: advertised and received
  Received 458 messages, 0 notifications, 0 in queue
  Sent 457 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 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: VPNv4 Unicast  BGP table version 4, neighbor version 4
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 5, neighbor version 5
  Index 1, Offset 0, Mask 0x2
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  3 accepted prefixes
  3 announced prefixes


For address family: VPNv4 Flow Specification  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 11.1.1.2, Local port: 179
Foreign host: 11.1.1.1, Foreign port: 41446
Nexthop: 11.1.1.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 03:09:46, due to Administratively Reset (Cease Notification sent)
Notification Error Message: (Cease/Administratively Reset.)

BGP neighbor is 20.1.1.1, remote AS 65000, local AS 65000, internal link
  BGP version 4, local router ID 2.2.2.2, remote router ID 1.1.1.1
  BGP state = Established, up for 03:10:17
  Last read 00:00:23, 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
    Address family VPNv4 Unicast: advertised and received
    Address family IPv4 Flow Specification: advertised and received
    Address family VPNv4 Flow Specification: advertised and received
  Received 457 messages, 1 notifications, 0 in queue
```

```
  Sent 457 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 2, Offset 0, Mask 0x4
  AIGP is enabled
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: VPNv4 Unicast  BGP table version 4, neighbor version 4
  Index 2, Offset 0, Mask 0x4
  AIGP is enabled
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 5, neighbor version 5
  Index 2, Offset 0, Mask 0x4
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  3 accepted prefixes
  3 announced prefixes


For address family: VPNv4 Flow Specification  BGP table version 3, neighbor version 3
  Index 2, Offset 0, Mask 0x4
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 20.1.1.2, Local port: 43160
Foreign host: 20.1.1.1, Foreign port: 179
Nexthop: 20.1.1.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 03:10:22, due to BGP Notification received
Notification Error Message: (OPEN Message Error/Bad BGP Identifier. Reason: )

RR2#
RR2#
RR2#sh ip bgp ipv4 flowspec summary
BGP router identifier 2.2.2.2, local AS number 65000
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
11.1.1.1                 4 65000  458        459       5      0      0  03:09:54              3
20.1.1.1                 4 65000  459        457       5      0      0  03:10:25              3

Total number of neighbors 2

Total number of Established sessions 2
RR2#
RR2#sh ip bgp ipv4 flowspec
```

```
BGP table version is 5, local router ID is 2.2.2.2
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
*>i  flowspec/216     20.1.1.1              0        100       0   100 i
*>i  flowspec/336     11.1.1.1              0        100       0   200 i
*>i  flowspec/264     11.1.1.1              0        100       0    i
*>i  flowspec/392     20.1.1.1              0        100       0    i
*>i  flowspec/176     20.1.1.1              0        100       0   100 i
*>i  flowspec/288     11.1.1.1              0        100       0   200 i

Total number of prefixes 6
RR2#
RR2#
RR2#sh ip bgp ipv4 flowspec detail
BGP table version is 5, local router ID is 2.2.2.2
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
*>i  flowspec/216     20.1.1.1              0        100       0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:44
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>i  flowspec/336     11.1.1.1              0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:11
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:100
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:200
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*>i  flowspec/264     11.1.1.1              0        100       0    i
```

```
  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:111
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:230
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:299

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*>i  flowspec/392     20.1.1.1             0        100       0    i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 1.1.1.1/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:78
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:101
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:201
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:400
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:499
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:330
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    Packet length:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x10

  FlowSpec Actions:
    traffic_rate: 1000011 bps
    traffic marking DSCP: e

*>i  flowspec/176     20.1.1.1             0        100       0   100 i

  FlowSpec:
    destination ip address: 4.4.4.4/32
    source ip address: 4.4.4.4/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:12
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:333
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:444

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>i  flowspec/288     11.1.1.1             0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:121
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:777
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:888
```

© 2026 IP Infusion Inc. Proprietary

```
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1001
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1200

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 6
RR2#
RR2#
RR2#sh ip bgp ipv4 flowspec vrf all
BGP table version is 5, local router ID is 2.2.2.2
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
*>i  flowspec/216     20.1.1.1              0        100       0   100 i
*>i  flowspec/336     11.1.1.1              0        100       0   200 i
*>i  flowspec/264     11.1.1.1              0        100       0    i
*>i  flowspec/392     20.1.1.1              0        100       0    i
*>i  flowspec/176     20.1.1.1              0        100       0   100 i
*>i  flowspec/288     11.1.1.1              0        100       0   200 i

Total number of prefixes 6
RR2#
RR2#sh ip bgp ipv4 flowspec vrf all detail
BGP table version is 5, local router ID is 2.2.2.2
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
*>i  flowspec/216     20.1.1.1              0        100       0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:44
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>i  flowspec/336     11.1.1.1              0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:11
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:100
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:200
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:400
    TCP/UDP source port:
```

© 2026 IP Infusion Inc. Proprietary

```
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*>i  flowspec/264     11.1.1.1             0        100       0    i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:111
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:230
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:299

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*>i  flowspec/392     20.1.1.1             0        100       0    i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 1.1.1.1/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:78
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:101
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:201
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:400
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:499
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:330
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    Packet length:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x10

  FlowSpec Actions:
    traffic_rate: 1000011 bps
    traffic marking DSCP: e

*>i  flowspec/176     20.1.1.1             0        100       0   100 i

  FlowSpec:
    destination ip address: 4.4.4.4/32
    source ip address: 4.4.4.4/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:12
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:333
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:444
```

```
  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>i  flowspec/288     11.1.1.1             0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:121
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:777
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:888
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1001
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1200

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 6
RR2#
RR2#
RR2#sh ip bgp vpnv4 all
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
Route Distinguisher: 100:1
*>i  4.4.4.4/32       20.1.1.1             0        100       0   100 ?
*>i  5.5.5.5/32       11.1.1.1             0        100       0   200 ?
*>i  10.1.1.0/24      20.1.1.1             0        100       0   100 ?
*>i  30.1.1.0/24      11.1.1.1             0        100       0   200 ?
*>i  55.1.1.0/24      11.1.1.1             0        100       0    ?
*>i  66.1.1.0/24      20.1.1.1             0        100       0    ?
*>i  100.1.1.0/24     20.1.1.1             0        100       0   100 ?
*>i  200.1.1.0        11.1.1.1             0        100       0   200 ?
 Announced routes count = 0
 Accepted routes count = 8
RR2#
RR2#sh ip bgp vpnv4 flowspec summary
BGP router identifier 2.2.2.2, local AS number 65000
BGP table version is 3
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
11.1.1.1                 4 65000  460        460       3      0      0  03:10:38              2
20.1.1.1                 4 65000  460        459       3      0      0  03:11:09              2

Total number of neighbors 2

Total number of Established sessions 2
RR2#
```

PE2

```
PE2#sh 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
RR2            ce43        3417.ebe4.af2c  Up     27        L1   IS-IS
                                           Up     27        L2   IS-IS
RR1            ce4         5c07.5871.205c  Up     8         L1   IS-IS
                                           Up     8         L2   IS-IS
PE2#
PE2#sh ip bgp
BGP table version is 4, local router ID is 3.3.3.3
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
*>i  4.4.4.4/32       1.1.1.1               0        100       0   100 ?
* i                   20.1.1.1              0        100       0   100 ?
*>   5.5.5.5/32       30.1.1.1              0        100       0   200 ?
*>i  10.1.1.0/24      1.1.1.1               0        100       0   100 ?
* i                   20.1.1.1              0        100       0   100 ?
*>   30.1.1.0/24      30.1.1.1              0        100       0   200 ?
*>   55.1.1.0/24      30.1.1.1              0        100       0   200 ?
*>i  66.1.1.0/24      1.1.1.1               0        100       0   100 ?
* i                   20.1.1.1              0        100       0   100 ?
*>i  100.1.1.0/24     1.1.1.1               0        100       0   100 ?
* i                   20.1.1.1              0        100       0   100 ?
*>   200.1.1.0        30.1.1.1              0        100       0   200 ?

Total number of prefixes 8
PE2#
PE2#sh ip bgp neighbors
BGP neighbor is 6.6.6.6, remote AS 65000, local AS 65000, internal link
  BGP version 4, local router ID 3.3.3.3, remote router ID 6.6.6.6
  BGP state = Established, up for 03:12:03
  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
    Address family VPNv4 Unicast: advertised and received
    Address family IPv4 Flow Specification: advertised and received
    Address family VPNv4 Flow Specification: advertised and received
  Received 459 messages, 0 notifications, 0 in queue
  Sent 459 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 4, neighbor version 4
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: VPNv4 Unicast  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
```

```
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 6, neighbor version 6
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  3 accepted prefixes
  3 announced prefixes


For address family: VPNv4 Flow Specification  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 3.3.3.3, Local port: 179
Foreign host: 6.6.6.6, Foreign port: 35405
Nexthop: 3.3.3.3
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 11.1.1.2, remote AS 65000, local AS 65000, internal link
  BGP version 4, local router ID 3.3.3.3, remote router ID 2.2.2.2
  BGP state = Established, up for 03:12:52
  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
    Address family VPNv4 Unicast: advertised and received
    Address family IPv4 Flow Specification: advertised and received
    Address family VPNv4 Flow Specification: advertised and received
  Received 462 messages, 1 notifications, 0 in queue
  Sent 467 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 4, neighbor version 4
  Index 2, Offset 0, Mask 0x4
  AIGP is enabled
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: VPNv4 Unicast  BGP table version 3, neighbor version 3
  Index 2, Offset 0, Mask 0x4
  AIGP is enabled
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 6, neighbor version 6
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  3 accepted prefixes
```

```
  3 announced prefixes


For address family: VPNv4 Flow Specification  BGP table version 3, neighbor version 3
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 11.1.1.1, Local port: 41446
Foreign host: 11.1.1.2, Foreign port: 179
Nexthop: 11.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 03:12:57, due to BGP Notification received
Notification Error Message: (OPEN Message Error/Bad BGP Identifier. Reason: )

BGP neighbor is 30.1.1.1, remote AS 200, local AS 65000, external link
  BGP version 4, local router ID 3.3.3.3, remote router ID 5.5.5.5
  BGP state = Established, up for 03:12:27
  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
    Address family IPv4 Flow Specification: advertised and received
  Received 456 messages, 0 notifications, 0 in queue
  Sent 463 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 3, Offset 0, Mask 0x8
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  4 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 6, neighbor version 6
  Index 3, Offset 0, Mask 0x8
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  4 announced prefixes

 Connections established 1; dropped 0
Local host: 30.1.1.2, Local port: 179
Foreign host: 30.1.1.1, Foreign port: 41276
Nexthop: 30.1.1.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 03:12:27, due to OPEN Message Error (Notification sent)
Notification Error Message: (OPEN Message Error/Bad BGP Identifier. Reason: )

BGP neighbor is 55.1.1.2, vrf vrf100, remote AS 200, local AS 65000, external link
  BGP version 4, local router ID 55.1.1.1, remote router ID 5.5.5.5
  BGP state = Established, up for 03:12:27
  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
    Address family IPv4 Flow Specification: advertised and received
  Received 460 messages, 0 notifications, 0 in queue
  Sent 460 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 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (standard)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  5 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 55.1.1.1, Local port: 40660
Foreign host: 55.1.1.2, Foreign port: 179
Nexthop: 55.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 03:12:32, due to OPEN Message Error (Notification sent)
Notification Error Message: (OPEN Message Error/Bad BGP Identifier. Reason: )

PE2#
PE2#sh ip bgp ipv4 flowspec summary
BGP router identifier 3.3.3.3, local AS number 65000
BGP table version is 6
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
6.6.6.6                  4 65000  460        459       6      0      0  03:12:13               3
11.1.1.2                 4 65000  464        468       6      0      0  03:13:02               3
30.1.1.1                 4   200  456        464       6      0      0  03:12:37               2

Total number of neighbors 3

Total number of Established sessions 3
BGP router identifier 55.1.1.1, local AS number 65000
BGP VRF vrf100 Route Distinguisher: 100:1
BGP table version is 1
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
55.1.1.2                 4   200  460        461       1      0      0  03:12:37               2

Total number of neighbors 1

Total number of Established sessions 1
PE2#
PE2#sh ip bgp ipv4 flowspec
BGP table version is 6, local router ID is 3.3.3.3
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
*>i  flowspec/216     6.6.6.6               0        100       0   100 i
* i                   11.1.1.2              0        100       0   100 i
*>   flowspec/336     30.1.1.1              0        100       0   200 i
*>   flowspec/264     0.0.0.0               0        100       32768  i
*>i  flowspec/392     6.6.6.6               0        100       0    i
```

```
* i                   11.1.1.2             0        100       0    i
*>i  flowspec/176     6.6.6.6              0        100       0   100 i
* i                   11.1.1.2             0        100       0   100 i
*>   flowspec/288     30.1.1.1             0        100       0   200 i

Total number of prefixes 6
PE2#
PE2#sh ip bgp ipv4 flowspec detail
BGP table version is 6, local router ID is 3.3.3.3
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
*>i  flowspec/216     6.6.6.6               0        100       0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:44
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

* i                   11.1.1.2             0        100       0   100 i
  FlowSpec Actions:
    traffic_rate: 21000 bps

*>   flowspec/336     30.1.1.1             0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:11
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:100
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:200
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*>   flowspec/264     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:111
```

```
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:230
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:299

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*>i  flowspec/392     6.6.6.6              0        100       0    i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 1.1.1.1/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:78
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:101
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:201
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:400
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:499
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:330
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    Packet length:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x10

  FlowSpec Actions:
    traffic_rate: 1000011 bps
    traffic marking DSCP: e

* i                   11.1.1.2             0        100       0    i
  FlowSpec Actions:
    traffic_rate: 1000011 bps
    traffic marking DSCP: e

*>i  flowspec/176     6.6.6.6              0        100       0   100 i

  FlowSpec:
    destination ip address: 4.4.4.4/32
    source ip address: 4.4.4.4/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:12
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:333
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:444

  FlowSpec Actions:
    traffic_rate: 4444444 bps

* i                   11.1.1.2             0        100       0   100 i
  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>   flowspec/288     30.1.1.1             0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    IP protocol:
```

```
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:121
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:777
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:888
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1001
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1200

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 6
PE2#
PE2#
PE2#sh ip bgp ipv4 flowspec vrf all
BGP table version is 1, local router ID is 55.1.1.1
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

BGP Route Table for VRF vrf100
*>i  flowspec/216     6.6.6.6               0        100       0   100 i
*>   flowspec/336     55.1.1.2              0        100       0   200 i
*>i  flowspec/176     6.6.6.6               0        100       0   100 i
*>   flowspec/288     55.1.1.2              0        100       0   200 i

Total number of prefixes 4
*>i  flowspec/216     6.6.6.6               0        100       0   100 i
* i                   11.1.1.2              0        100       0   100 i
*>   flowspec/336     30.1.1.1              0        100       0   200 i
*>   flowspec/264     0.0.0.0               0        100       32768  i
*>i  flowspec/392     6.6.6.6               0        100       0    i
* i                   11.1.1.2              0        100       0    i
*>i  flowspec/176     6.6.6.6               0        100       0   100 i
* i                   11.1.1.2              0        100       0   100 i
*>   flowspec/288     30.1.1.1              0        100       0   200 i

Total number of prefixes 6
PE2#
PE2#sh ip bgp ipv4 flowspec vrf all detail
BGP table version is 1, local router ID is 55.1.1.1
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

BGP Route Table for VRF vrf100
*>i  flowspec/216     6.6.6.6               0        100       0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:44
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122
```

```
  FlowSpec Actions:
    traffic_rate: 21000 bps

*>   flowspec/336     55.1.1.2             0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:11
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:100
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:200
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*>i  flowspec/176     6.6.6.6              0        100       0   100 i

  FlowSpec:
    destination ip address: 4.4.4.4/32
    source ip address: 4.4.4.4/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:12
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:333
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:444

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>   flowspec/288     55.1.1.2             0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:121
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:777
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:888
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1001
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1200

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 4
*>i  flowspec/216     6.6.6.6              0        100       0   100 i

  FlowSpec:
```

```
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:44
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

* i                   11.1.1.2             0        100       0   100 i
  FlowSpec Actions:
    traffic_rate: 21000 bps

*>   flowspec/336     30.1.1.1             0        100       0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:11
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:100
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:200
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*>   flowspec/264     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:111
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:230
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:299

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*>i  flowspec/392     6.6.6.6              0        100       0    i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 1.1.1.1/32
    IP protocol:
```

```
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:78
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:101
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:201
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:400
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:499
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:330
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    Packet length:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x10

  FlowSpec Actions:
    traffic_rate: 1000011 bps
    traffic marking DSCP: e

* i                   11.1.1.2             0        100       0    i
  FlowSpec Actions:
    traffic_rate: 1000011 bps
    traffic marking DSCP: e

*>i  flowspec/176     6.6.6.6              0        100       0   100 i

  FlowSpec:
    destination ip address: 4.4.4.4/32
    source ip address: 4.4.4.4/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:12
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:333
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:444

  FlowSpec Actions:
    traffic_rate: 4444444 bps

* i                   11.1.1.2             0        100       0   100 i
  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>   flowspec/288     30.1.1.1             0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:121
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:777
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:888
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1001
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1200

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 6
PE2#
```

```
PE2#sh ip bgp vpnv4 flowspec summary
BGP router identifier 3.3.3.3, local AS number 65000
BGP table version is 3
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
6.6.6.6                  4 65000  463        463       3      0      0  03:13:42               2
11.1.1.2                 4 65000  467        471       3      0      0  03:14:31               2

Total number of neighbors 2

Total number of Established sessions 2
PE2#
PE2#sh ip bgp vpnv4 flowspec vrf vrf100
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
Route Distinguisher: 100:1 (Default for VRF vrf100)
*>i  flowspec/216     6.6.6.6               0        100        0   100 i
*>   flowspec/336     55.1.1.2              0        100        0   200 i
*>i  flowspec/176     6.6.6.6               0        100        0   100 i
*>   flowspec/288     55.1.1.2              0        100        0   200 i
 Announced routes count = 2
 Accepted routes count = 2
PE2#
PE2#sh ip bgp vpnv4 flowspec vrf vrf100 detail
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
Route Distinguisher: 100:1 (Default for VRF vrf100)
*>i  flowspec/216     6.6.6.6               0        100        0   100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:44
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*>   flowspec/336     55.1.1.2              0        100        0   200 i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:11
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:100
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:200
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:400
```

© 2026 IP Infusion Inc. Proprietary

```
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*>i  flowspec/176     6.6.6.6              0        100       0   100 i

  FlowSpec:
    destination ip address: 4.4.4.4/32
    source ip address: 4.4.4.4/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:12
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:333
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:444

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>   flowspec/288     55.1.1.2             0        100       0   200 i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:121
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:777
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:888
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1001
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1200

  FlowSpec Actions:
    traffic_rate: 100000 bps

 Announced routes count = 2
 Accepted routes count = 2
PE2#
```

CE2

```
CE2#sh ip bgp
BGP table version is 3, local router ID is 5.5.5.5
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
*>   4.4.4.4/32       30.1.1.2              0        100       0   65000 100 ?
*                     55.1.1.1              0        100       0   65000 100 ?
*>   5.5.5.5/32       0.0.0.0               0        100       32768  ?
*>   10.1.1.0/24      30.1.1.2              0        100       0   65000 100 ?
*                     55.1.1.1              0        100       0   65000 100 ?
*>   30.1.1.0/24      0.0.0.0               0        100       32768  ?
*>   55.1.1.0/24      0.0.0.0               0        100       32768  ?
*                     55.1.1.1              0        100       0   65000 ?
```

```
*>   66.1.1.0/24      55.1.1.1            0        100       0   65000 ?
*                     30.1.1.2            0        100       0   65000 100 ?
*>   100.1.1.0/24     30.1.1.2            0        100       0   65000 100 ?
*                     55.1.1.1            0        100       0   65000 100 ?
*>   200.1.1.0        0.0.0.0             0        100       32768  ?

Total number of prefixes 8
CE2#
CE2#
CE2#sh ip bgp neighbors
BGP neighbor is 30.1.1.2, remote AS 65000, local AS 200, external link
  BGP version 4, local router ID 5.5.5.5, remote router ID 3.3.3.3
  BGP state = Established, up for 03:15:13
  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 IPv4 Unicast: advertised and received
    Address family IPv4 Flow Specification: advertised and received
  Received 467 messages, 1 notifications, 0 in queue
  Sent 463 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)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  5 announced prefixes


For address family: IPv4 Flow Specification  BGP table version 2, neighbor version 2
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 30.1.1.1, Local port: 41276
Foreign host: 30.1.1.2, Foreign port: 179
Nexthop: 30.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 03:15:18, due to BGP Notification received
Notification Error Message: (OPEN Message Error/Bad BGP Identifier. Reason: )

BGP neighbor is 55.1.1.1, remote AS 65000, local AS 200, external link
  BGP version 4, local router ID 5.5.5.5, remote router ID 55.1.1.1
  BGP state = Established, up for 03:15:13
  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
    Address family IPv4 Flow Specification: advertised and received
  Received 466 messages, 1 notifications, 0 in queue
  Sent 467 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
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  5 accepted prefixes
  7 announced prefixes
```

```
For address family: IPv4 Flow Specification  BGP table version 2, neighbor version 2
  Index 2, Offset 0, Mask 0x4
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  2 accepted prefixes
  2 announced prefixes

 Connections established 1; dropped 0
Local host: 55.1.1.2, Local port: 179
Foreign host: 55.1.1.1, Foreign port: 40660
Nexthop: 55.1.1.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 03:15:18, due to BGP Notification received
Notification Error Message: (OPEN Message Error/Bad BGP Identifier. Reason: )

CE2#
CE2#
CE2#sh ip bgp ipv4 flowspec summary
BGP router identifier 5.5.5.5, local AS number 200
BGP table version is 2
3 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/PfxRcd
30.1.1.2                 4 65000  468        463       2      0      0  03:15:21               4
55.1.1.1                 4 65000  467        467       2      0      0  03:15:21               2

Total number of neighbors 2

Total number of Established sessions 2
CE2#
CE2#
CE2#sh ip bgp ipv4 flowspec
BGP table version is 2, local router ID is 5.5.5.5
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
*    flowspec/216     30.1.1.2              0        100       0   65000 100 i
*                     55.1.1.1              0        100       0   65000 100 i
*>   flowspec/336     0.0.0.0               0        100       32768  i
*    flowspec/264     30.1.1.2              0        100       0   65000 i
*    flowspec/392     30.1.1.2              0        100       0   65000 i
*    flowspec/176     30.1.1.2              0        100       0   65000 100 i
*                     55.1.1.1              0        100       0   65000 100 i
*>   flowspec/288     0.0.0.0               0        100       32768  i

Total number of prefixes 6
CE2#
CE2#sh ip bgp ipv4 flowspec detail
BGP table version is 2, local router ID is 5.5.5.5
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
*    flowspec/216     30.1.1.2              0        100       0   65000 100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:44
    TCP/UDP destination port:
```

```
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*                     55.1.1.1             0        100       0   65000 100 i
  FlowSpec Actions:
    traffic_rate: 21000 bps

*>   flowspec/336     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:11
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:100
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:200
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*    flowspec/264     30.1.1.2             0        100       0   65000 i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:111
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:230
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:299

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*    flowspec/392     30.1.1.2             0        100       0   65000 i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 1.1.1.1/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:78
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:101
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:201
    TCP/UDP destination port:
```

```
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:400
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:499
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:330
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    Packet length:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x10

  FlowSpec Actions:
    traffic_rate: 1000011 bps
    traffic marking DSCP: e

*    flowspec/176     30.1.1.2             0        100       0   65000 100 i

  FlowSpec:
    destination ip address: 4.4.4.4/32
    source ip address: 4.4.4.4/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:12
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:333
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:444

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*                    55.1.1.1             0        100       0   65000 100 i
  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>   flowspec/288     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:121
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:777
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:888
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1001
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1200

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 6
CE2#
CE2#
CE2#sh ip bgp ipv4 flowspec vrf all
BGP table version is 2, local router ID is 5.5.5.5
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
*    flowspec/216     30.1.1.2              0        100       0   65000 100 i
*                     55.1.1.1              0        100       0   65000 100 i
```

```
*>   flowspec/336     0.0.0.0              0        100       32768  i
*    flowspec/264     30.1.1.2             0        100       0   65000 i
*    flowspec/392     30.1.1.2             0        100       0   65000 i
*    flowspec/176     30.1.1.2             0        100       0   65000 100 i
*                     55.1.1.1             0        100       0   65000 100 i
*>   flowspec/288     0.0.0.0              0        100       32768  i

Total number of prefixes 6
CE2#
CE2#sh ip bgp ipv4 flowspec vrf all detail
BGP table version is 2, local router ID is 5.5.5.5
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
*    flowspec/216     30.1.1.2              0        100       0   65000 100 i

  FlowSpec:
    destination ip address: 10.1.1.0/24
    source ip address: 10.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:44
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:44
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:144
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:122

  FlowSpec Actions:
    traffic_rate: 21000 bps

*                    55.1.1.1             0        100       0   65000 100 i
  FlowSpec Actions:
    traffic_rate: 21000 bps

*>   flowspec/336     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 30.1.1.0/24
    source ip address: 30.1.1.0/24
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:11
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:100
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:200
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:500
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:600
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x21
      e:1 a:1 len:0 not:1 m:0 data:0x1e
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x21

  FlowSpec Actions:
    traffic_rate: 500000 bps
    traffic marking DSCP: 1e

*    flowspec/264     30.1.1.2             0        100       0   65000 i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 3.3.3.3/32
```

```
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:111
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:230
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:299

  FlowSpec Actions:
    traffic_rate: 1111111 bps

*    flowspec/392     30.1.1.2             0        100       0   65000 i

  FlowSpec:
    destination ip address: 3.3.3.3/32
    source ip address: 1.1.1.1/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:78
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:101
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:201
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:400
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:499
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:330
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:399
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    Packet length:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:22
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:222
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x10

  FlowSpec Actions:
    traffic_rate: 1000011 bps
    traffic marking DSCP: e

*    flowspec/176     30.1.1.2             0        100       0   65000 100 i

  FlowSpec:
    destination ip address: 4.4.4.4/32
    source ip address: 4.4.4.4/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:12
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:333
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:444

  FlowSpec Actions:
    traffic_rate: 4444444 bps

*                    55.1.1.1             0        100       0   65000 100 i
  FlowSpec Actions:
    traffic_rate: 4444444 bps

*>   flowspec/288     0.0.0.0              0        100       32768  i

  FlowSpec:
    destination ip address: 5.5.5.5/32
    source ip address: 5.5.5.5/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:121
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:777
```

```
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:888
    TCP/UDP destination port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1300
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1400
    TCP/UDP source port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1001
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:1200

  FlowSpec Actions:
    traffic_rate: 100000 bps


Total number of prefixes 6
CE2#
```

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 90. Three-way handshake

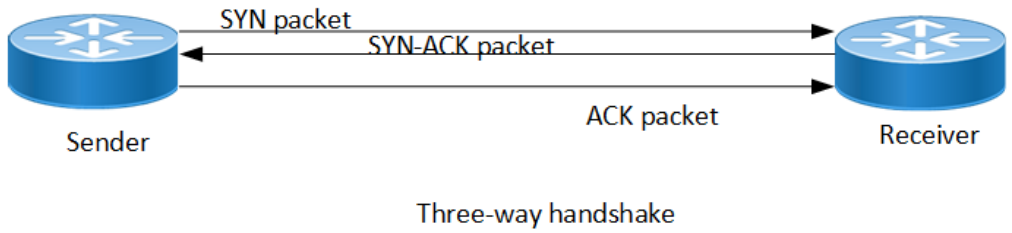

Feature Characteristics

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 91. 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.

© 2026 IP Infusion Inc. Proprietary

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 92. 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`	Define BGP neighbors, and establish a TCP session.

© 2026 IP Infusion Inc. Proprietary

`remoteas 100`	`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.

© 2026 IP Infusion Inc. Proprietary

`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)
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) command. When you specify this parameter, the command applies to all peers in the group.

<40-1440>
Configure TCP MSS

Default

Disabled

Command Mode

Router mode, address family-vrf mode and BGP unnumbered mode.

Applicability

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

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

BGP ORF Prefix-List VPNV4 Address

Overview

The Border Gateway Protocol (BGP) Outbound Route Filtering feature operates as a Prefix-Based filtering system within the BGP. Its primary purpose is to reduce the volume of BGP updates exchanged among peer routers. By selectively screening out unnecessary routing updates at the source, this feature effectively lessens the strain on resources needed for generating and handling routing updates. Its objective is to streamline router processing, especially for routers not set up to accept full BGP route updates from a service provider network.

Feature Characteristics

This feature provides customers with various routing options, such as access to routing information like a full table view, solely a default route, or a tailored subset such as a default route combined with locally originated prefixes from the service provider. Typically, BGP service providers do not impose complex outbound filtering policies on their customers.

Benefits

The advantages of Prefix-Based Outbound Route Filtering:

- Minimize unnecessary routing updates
- Reduces resources required for routing update generation and processing
- Reduces required to receive and discard routes.

Configuration

The BGP Prefix-Based Outbound Route Filtering feature offers support for prefix length matching, wildcard-based prefix matching, and exact address prefix matching across various address families. It allows configuration on a router to enable Outbound Route Filtering (ORF) capabilities for sending or receiving, using the "send" or "receive" keywords. Moreover, it permits configuration to enable both sending and receiving ORF capabilities using the "both" keyword.

Topology

In this topology, the PE1, P1, RR1, and PE2 interface is established. It allows configuration on a router to enable Outbound Route Filtering (ORF) capabilities.

© 2026 IP Infusion Inc. Proprietary

Figure 93. ORF-Prefix VPNv4 Address Topology

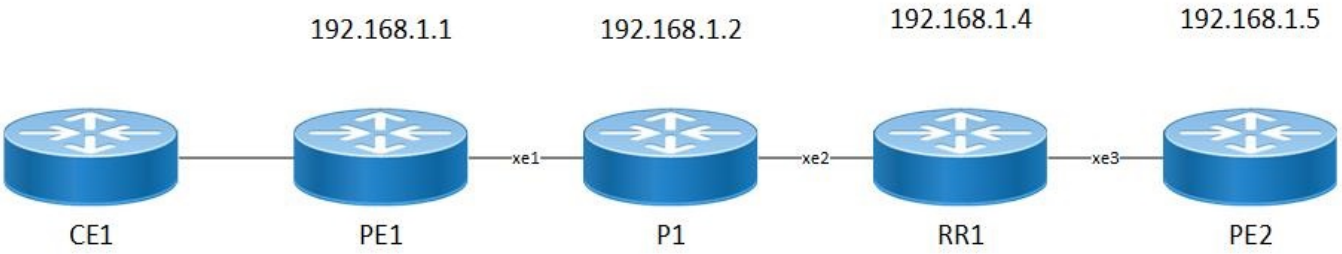

PE1

The Provider Edge (PE) 1 is a device at the edge of a service provider's network. For the PE1 configuration, follow these steps.

1. Enable activate the loopback interface, enter the following command while in configuration mode. Then, proceed to set up the IP address for the loopback interface.

```
PE1(config)#interface lo
PE1(config-if)#ip address 192.168.1.1/32 secondary
PE1(config-if)#enable-ldp ipv4
```

2. To exit from the loopback interface, execute the following command.

```
PE1(config)#exit
```

3. Set up the Router ID. Next, configure targeted LDP sessions for PE-1. Once done, exit from targeted-peer mode. Then, configure the transport address for LDP to run on an IPv4 interface for TCP sessions.

```
PE1(config)#router ldp
PE1(config-router)#router-id 192.168.1.1
PE1(config-router)#targeted-peer ipv4 192.168.1.5
PE1(config-router-targeted-peer)#exit-targeted-peer-mode
PE1(config-router)#transport-address ipv4 192.168.1.1
```

4. To exit from the router mode for LDP, execute the following command.

```
PE1(config)#exit
```

5. In interface mode, assign an IP address to the interface. Then, activate label switching capability on the interface and enable LDP on it.

```
PE1(config)#interface xe1
PE1(config-if)#ip address 12.1.1.1/24
PE1(config-if)#label-switching
PE1(config-if)#enable-ldp ipv4
PE1(config-if)# ip ospf cost 10
```

6. To exit from the interface configuration on network side, execute the following command.

```
PE1(config)#exit
```

7. Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process.

```
PE1(config)# router ospf 100
```

© 2026 IP Infusion Inc. Proprietary

8. Configure the OSPF Router ID and define the interface for OSPF operation and link it with the area ID (0).

```
PE1(config-router)# ospf router-id 192.168.1.1
PE1(config-router)# bfd all-interfaces
PE1(config-router)# timers spf exp 50 50
PE1(config-router)# timers throttle lsa all 0 1 1
PE1(config-router)# network 12.1.1.0/24 area 0.0.0.0
PE1(config-router)# network 192.168.1.1/32 area 0.0.0.0
```

9. To exit from the OSPF, execute the following command.

```
PE1(config)#exit
```

10. Switch to BGP router mode. Establish PE1 as an iBGP peer. Specify the loopback as the source for iBGP peering with the remote PE1 router. Activate PE1 in the VPNv4 unicast address family.

```
PE1(config)#router bgp 100
PE1(config-router)# bgp router-id 192.168.1.1
PE1(config-router)# neighbor 192.168.1.4 remote-as 100
PE1(config-router)# neighbor 192.168.1.4 update-source lo
PE1(config-router)# neighbor 192.168.1.4 advertisement-interval 0
PE1(config-router)# address-family vpnv4 unicast
PE1(config-router-af)# neighbor 192.168.1.4 activate
PE1(config-router-af)# neighbor 192.168.1.4 capability orf prefix-list receive
PE1(config-router-af)# exit-address-family
```

11. To exit from the BGP, execute the following command.

```
PE1(config)#exit
```

P

The Provider (P) is a device at the edge of a service provider's network. For the P configuration, follow these steps.

1. Enable activate the loopback interface, enter the following command while in configuration mode. Then, proceed to set up the IP address for the loopback interface.

```
P(config)#interface lo
P(config-if)#ip address 192.168.1.1/32 secondary
P(config-if)#enable-ldp ipv4
```

2. To exit from the loopback interface, execute the following command.

```
P(config)#exit
```

3. Set up the Router ID. Next, configure targeted LDP sessions for P. Once done, exit from targeted-peer mode. Then, configure the transport address for LDP to run on an IPv4 interface for TCP sessions.

```
P(config)#router ldp
P(config-router)#router-id 192.168.1.2
P(config-router)#targeted-peer ipv4 192.168.1.2
```

4. To exit from the router mode for LDP, execute the following command.

```
P(config)#exit
```

5. In interface mode, assign an IP address to the interface. Then, activate label switching capability on the interface and enable LDP on it.

```
P(config)#interface xe1
P(config-if)#ip address 12.1.1.1/24
P(config-if)#label-switching
P(config)#interface xe1
P(config-if)#ip address 13.1.1.2/24
P(config-if)#label-switching
```

6. To exit from the interface configuration on network side, execute the following command.

```
P(config)#exit
```

7. Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process.

```
P(config)# router ospf 100
```

8. Configure the OSPF Router ID and define the interface for OSPF operation and link it with the area ID (0).

```
P(config-router)# ospf router-id 192.168.1.2
P(config-router)# bfd all-interfaces
P(config-router)# timers spf exp 50 50
P(config-router)# timers throttle lsa all 0 1 1
P(config-router)# network 12.1.1.0/24 area 0.0.0.0
P(config-router)# network 192.168.1.1/32 area 0.0.0.0
```

9. To exit from the OSPF, execute the following command.

```
P(config)#exit
```

RR

Route Reflector (RR) is a designated router that will reflect routes learned from other iBGP peers. All routers form a peering relationship only with the Route Reflector. For the RR configuration, follow these steps.

1. Enable activate the loopback interface, enter the following command while in configuration mode. Then, proceed to set up the IP address for the loopback interface.

```
RR(config)#interface lo
RR(config-if)#ip address 192.168.1.1/32 secondary
RR(config-if)#enable-ldp ipv4
```

2. To exit from the loopback interface, execute the following command.

```
RR(config)#exit
```

3. Set up the Router ID. Next, configure targeted LDP sessions for P. Once done, exit from targeted-peer mode. Then, configure the transport address for LDP to run on an IPv4 interface for TCP sessions.

```
RR(config)#router ldp
RR(config-router)#router-id 192.168.1.4/32
RR(config-router)#transport-address ipv4 192.168.1.4
```

4. To exit from the router mode for LDP, execute the following command.

```
RR(config)#exit
```

5. In interface mode, set up the IP address for the interface and activate label switching capability on it.

```
RR(config)#interface xe1
RR(config-if)#ip address 12.1.1.1/24
RR(config-if)#label-switching
RR(config-if)#enable-ldp ipv4
RR(config)#interface xe3
RR(config-if)#ip address 14.1.1.1/24
RR(config-if)#label-switching
RR(config-if)#enable-ldp ipv4
```

6. To exit from the interface configuration on network side, execute the following command.

```
RR(config)#exit
```

7. Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process.

```
RR(config)# router ospf 100
```

© 2026 IP Infusion Inc. Proprietary

8. Configure the OSPF Router ID and define the interface for OSPF operation and link it with the area ID (0).

```
RR(config-router)# ospf router-id 192.168.1.4
RR(config-router)# bfd all-interfaces
RR(config-router)# timers spf exp 50 50
RR(config-router)# timers throttle lsa all 0 1 1
RR(config-router)# network 12.1.1.0/24 area 0.0.0.0
RR(config-router)# network 192.168.1.1/32 area 0.0.0.0
```

9. To exit from the OSPF, execute the following command.

```
RR(config)#exit
```

10. Switch to BGP router mode. Establish RR as an iBGP peer. Specify the loopback as the source for iBGP peering with the remote RR router. Activate RR in the VPNv4 unicast address family.

```
RR(config)#router bgp 100
RR(config-router)# bgp router-id 192.168.1.1
RR(config-router)# neighbor 192.168.1.4 remote-as 100
RR(config-router)# neighbor 192.168.1.4 update-source lo
RR(config-router)# neighbor 192.168.1.4 advertisement-interval 0
RR(config-router)# address-family vpnv4 unicast
RR(config-router-af)# neighbor 192.168.1.4 active
RR(config-router-af)# neighbor 192.168.1.4 capability orf prefix-list s
RR(config-router-af)# neighbor 192.168.1.1 prefix-list
RR(config-router-af)# exit-address-family
```

11. To exit from the BGP, execute the following command.

```
PE1(config)#exit
```

12. To configure the global prefix, execute the following command in the global mode.

```
RR(config)# ip prefix-list ORF1
RR(config-ip-prefix-list)# seq 1 permit 45.1.1.0/24
```

13. To exit from the BGP, execute the following command.

```
RR(config)#exit
```

PE2

The Provider Edge (PE) 2 is a device at the edge of a service provider's network. For the PE2 configuration, follow these steps.

1. Enable activate the loopback interface, enter the following command while in configuration mode. Then, proceed to set up the IP address for the loopback interface.

```
PE2(config)#interface lo
PE2(config-if)#ip address 192.168.1.1/32 secondary
PE2(config-if)#enable-ldp ipv4
```

2. To exit from the loopback interface, execute the following command.

```
PE2(config)#exit
```

3. Set up the Router ID. Next, configure targeted LDP sessions for P. Once done, exit from targeted-peer mode. Then, configure the transport address for LDP to run on an IPv4 interface for TCP sessions.

```
PE2(config)#router ldp
PE2(config-router)#router-id 192.168.1.1
PE2(config-router)#targeted-peer ipv4 192.168.1.5
PE2(config-router-targeted-peer)#exit-targeted-peer-mode
PE2(config-router)#transport-address ipv4 192.168.1.1
PE2(config-router)#transport-address ipv4 192.168.1.5
```

4. To exit from the router mode for LDP, execute the following command.

© 2026 IP Infusion Inc. Proprietary

```
PE2(config)#exit
```

5. In interface mode, assign an IP address to the interface. Then, activate label switching capability on the interface and enable LDP on it.

```
PE2(config)#interface xe1
PE2(config-if)#ip address 12.1.1.1/24
PE2(config-if)#label-switching
PE2(config-if)#enable-ldp ipv4
PE2(config-if)# ip ospf cost 10
```

6. To exit from the interface configuration on network side, execute the following command.

```
PE2(config)#exit
```

7. Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process.

```
PE2(config)# router ospf 100
```

8. Configure the OSPF Router ID and define the interface for OSPF operation and link it with the area ID (0).

```
PE2(config-router)# ospf router-id 192.168.1.1
PE2(config-router)# bfd all-interfaces
PE2(config-router)# timers spf exp 50 50
PE2(config-router)# timers throttle lsa all 0 1 1
PE2(config-router)# network 12.1.1.0/24 area 0.0.0.0
PE2(config-router)# network 192.168.1.1/32 area 0.0.0.0
```

9. To exit from the OSPF, execute the following command.

```
PE2(config)#exit
```

10. Switch to BGP router mode. Establish PE2 as an iBGP peer. Specify the loopback as the source for iBGP peering with the remote PE2 router. Activate PE2 in the VPNv4 unicast address family.

```
PE2(config)#router bgp 100
PE2(config-router)# bgp router-id 192.168.1.1
PE2(config-router)# neighbor 192.168.1.4 remote-as 100
PE2(config-router)# neighbor 192.168.1.4 update-source lo
PE2(config-router)# neighbor 192.168.1.4 advertisement-interval 0
PE2(config-router)# address-family vpnv4 unicast
PE2(config-router-af)# neighbor 192.168.1.4 activate
PE2(config-router-af)# exit-address-family
```

11. To exit from the BGP, execute the following command.

```
PE2(config)#exit
```

Validation

```
PE1-7017#show ip bgp vpnv4 all
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
Route Distinguisher: 192.168.1.1:100 (Default for VRF vrf100)
*> l 45.1.1.0/24      100.1.1.2             0        100        0   400 i
*> l 45.1.2.0/24      100.1.1.2             0        100        0   400 i
*> l 45.1.3.0/24      100.1.1.2             0        100        0   400 i
*> l 45.1.4.0/24      100.1.1.2             0        100        0   400 i
*> l 45.1.5.0/24      100.1.1.2             0        100        0   400 i
*> l 45.1.6.0/24      100.1.1.2             0        100        0   400 i
*> l 45.1.7.0/24      100.1.1.2             0        100        0   400 i
```

© 2026 IP Infusion Inc. Proprietary

```
*> l 45.1.8.0/24      100.1.1.2            0        100       0   400 i
*> l 45.1.9.0/24      100.1.1.2            0        100       0   400 i
*> l 45.1.10.0/24     100.1.1.2            0        100       0   400 i
*> l 100.1.1.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.2.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.3.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.4.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.5.0/24     0.0.0.0              0        100       32768  ?
*>il 200.1.1.0        192.168.1.4          0        100       0    ?
*>il 200.1.2.0        192.168.1.4          0        100       0    ?
*>il 200.1.3.0        192.168.1.4          0        100       0    ?
*>il 200.1.4.0        192.168.1.4          0        100       0    ?
*>il 200.1.5.0        192.168.1.4          0        100       0    ?
 Announced routes count = 15
 Accepted routes count = 5
Route Distinguisher: 192.168.1.1:101 (Default for VRF vrf101)
*>   45.1.1.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.2.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.3.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.4.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.5.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.6.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.7.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.8.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.9.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.10.0/24     100.1.1.2            0        100       0   400 i
*>   100.1.1.0/24     0.0.0.0              0        100       32768  ?
*> l 100.1.2.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.3.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.4.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.5.0/24     0.0.0.0              0        100       32768  ?
*>il 200.1.1.0        192.168.1.4          0        100       0    ?
*>il 200.1.2.0        192.168.1.4          0        100       0    ?
*>il 200.1.3.0        192.168.1.4          0        100       0    ?
*>il 200.1.4.0        192.168.1.4          0        100       0    ?
*>il 200.1.5.0        192.168.1.4          0        100       0    ?
 Announced routes count = 15
 Accepted routes count = 5
Route Distinguisher: 192.168.1.1:102 (Default for VRF vrf102)
*>   45.1.1.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.2.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.3.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.4.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.5.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.6.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.7.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.8.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.9.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.10.0/24     100.1.1.2            0        100       0   400 i
*>   100.1.1.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.2.0/24     0.0.0.0              0        100       32768  ?
*> l 100.1.3.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.4.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.5.0/24     0.0.0.0              0        100       32768  ?
*>il 200.1.1.0        192.168.1.4          0        100       0    ?
*>il 200.1.2.0        192.168.1.4          0        100       0    ?
*>il 200.1.3.0        192.168.1.4          0        100       0    ?
*>il 200.1.4.0        192.168.1.4          0        100       0    ?
*>il 200.1.5.0        192.168.1.4          0        100       0    ?
 Announced routes count = 15
 Accepted routes count = 5
Route Distinguisher: 192.168.1.1:103 (Default for VRF vrf103)
*>   45.1.1.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.2.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.3.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.4.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.5.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.6.0/24      100.1.1.2            0        100       0   400 i
```

```
*>   45.1.7.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.8.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.9.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.10.0/24     100.1.1.2            0        100       0   400 i
*>   100.1.1.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.2.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.3.0/24     0.0.0.0              0        100       32768  ?
*> l 100.1.4.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.5.0/24     0.0.0.0              0        100       32768  ?
*>il 200.1.1.0        192.168.1.4          0        100       0    ?
*>il 200.1.2.0        192.168.1.4          0        100       0    ?
*>il 200.1.3.0        192.168.1.4          0        100       0    ?
*>il 200.1.4.0        192.168.1.4          0        100       0    ?
*>il 200.1.5.0        192.168.1.4          0        100       0    ?
 Announced routes count = 15
 Accepted routes count = 5
Route Distinguisher: 192.168.1.1:104 (Default for VRF vrf104)
*>   45.1.1.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.2.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.3.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.4.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.5.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.6.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.7.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.8.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.9.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.10.0/24     100.1.1.2            0        100       0   400 i
*>   100.1.1.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.2.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.3.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.4.0/24     0.0.0.0              0        100       32768  ?
*> l 100.1.5.0/24     0.0.0.0              0        100       32768  ?
*>il 200.1.1.0        192.168.1.4          0        100       0    ?
*>il 200.1.2.0        192.168.1.4          0        100       0    ?
*>il 200.1.3.0        192.168.1.4          0        100       0    ?
*>il 200.1.4.0        192.168.1.4          0        100       0    ?
*>il 200.1.5.0        192.168.1.4          0        100       0    ?
 Announced routes count = 15
 Accepted routes count = 5
Route Distinguisher: 192.168.1.6:100
*>il 200.1.1.0        192.168.1.4          0        100       0    ?
* il                  192.168.1.5          0        100       0    ?
 Announced routes count = 0
 Accepted routes count = 2
Route Distinguisher: 192.168.1.6:101
*>il 200.1.2.0        192.168.1.4          0        100       0    ?
* il                  192.168.1.5          0        100       0    ?
 Announced routes count = 0
 Accepted routes count = 2
Route Distinguisher: 192.168.1.6:102
*>il 200.1.3.0        192.168.1.4          0        100       0    ?
* il                  192.168.1.5          0        100       0    ?
 Announced routes count = 0
 Accepted routes count = 2
Route Distinguisher: 192.168.1.6:103
*>il 200.1.4.0        192.168.1.4          0        100       0    ?
* il                  192.168.1.5          0        100       0    ?
 Announced routes count = 0
 Accepted routes count = 2
Route Distinguisher: 192.168.1.6:104
*>il 200.1.5.0        192.168.1.4          0        100       0    ?
* il                  192.168.1.5          0        100       0    ?
 Announced routes count = 0
 Accepted routes count = 2x
Use these commands to validate the BGP Neighbor Table.

PE1#show ip bgp neighbors 192.168.1.4
BGP neighbor is 192.168.1.4, remote AS 100, local AS 100, internal link, peer in
```

```
dex: 32
  BGP version 4, local router ID 192.168.1.1, remote router ID 192.168.1.4
  BGP state = Established, up for 00:05:03
  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 IPv4 Labeled-Unicast: advertised and received
    Address family VPNv4 Unicast: advertised and received
    Address family L2VPN VPLS: advertised and received
    Address family L2VPN EVPN: advertised and received
    Address family IPv6 Unicast: advertised and received
    Address family VPNv6 Unicast: advertised and received
    Address family IPv6 Labeled Unicast: advertised and received
  Received 3229 messages, 1 notifications, 0 in queue
  Sent 3252 messages, 2 notifications, 0 in queue
  Route refresh request: received 2, sent 0
  Minimum time between advertisement runs is 0 seconds
  Update source is lo

For address family: IPv4 Unicast  BGP table version 4, neighbor version 4
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  1 announced prefixes


For address family: VPNv4 Unicast  BGP table version 2, neighbor version 2
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  AF-dependant capabilities:
    Outbound Route Filter (ORF) type (64) Prefix-list:
      Send-mode: received
      Receive-mode: advertised
    Outbound Route Filter (ORF) type (128) Prefix-list:
      Send-mode: received
      Receive-mode: advertised
  Outbound Route Filter (ORF): received (1 entries)
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  5 accepted prefixes
  1 announced prefixes

For address family: IPv4 Labeled-Unicast  BGP table version 6, neighbor version
5
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  1 announced prefixes

For address family: L2VPN VPLS  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  9 announced prefixes

For address family: L2VPN EVPN  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  20 accepted prefixes
  Accepted AD:10 MACIP:0 MCAST:10 ESI:0 PREFIX:0
  20 announced prefixes
```

```
For address family: IPv6 Unicast  BGP table version 2, 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

For address family: VPNv6 Unicast  BGP table version 2, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  5 announced prefixes

For address family: 6PE Labeled Unicast  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  10 accepted prefixes
  10 announced prefixes

 Connections established 4; dropped 3
Local host: 192.168.1.1, Local port: 179
Foreign host: 192.168.1.4, Foreign port: 44897
TCP MSS: (0), Advertise TCP MSS: (9176), Send TCP MSS: (1460),  Receive TCP MSS:
 (1460)
Sock FD : (43)
Nexthop: 192.168.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:05:03, due to Administratively Reset (Cease Notification sent)
Notification Error Message: (Cease/Administratively Reset.)
Use these commands to validate the RR configuration.
# RR1-7038#SH IP BGP VPNV4 ALL
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
Route Distinguisher: 192.168.1.1:100
*>il 45.1.1.0/24      192.168.1.1           0        100        0   400 i
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 192.168.1.6:100
*>il 200.1.1.0        192.168.1.6           0        100        0    ?
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 192.168.1.6:101
*>il 200.1.2.0        192.168.1.6           0        100        0    ?
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 192.168.1.6:102
*>il 200.1.3.0        192.168.1.6           0        100        0    ?
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 192.168.1.6:103
*>il 200.1.4.0        192.168.1.6           0        100        0    ?
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 192.168.1.6:104
*>il 200.1.5.0        192.168.1.6           0        100        0    ?
 Announced routes count = 0
 Accepted routes count = 1
```

```
RR1-7038#show ip bgp neighbors 192.168.1.1
BGP neighbor is 192.168.1.1, remote AS 100, local AS 100, internal link, peer in
dex: 2
  BGP version 4, local router ID 192.168.1.4, remote router ID 192.168.1.1
  BGP state = Established, up for 00:06:19
  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
    Address family IPv4 Labeled-Unicast: advertised and received
    Address family VPNv4 Unicast: advertised and received
    Address family L2VPN VPLS: advertised and received
    Address family L2VPN EVPN: advertised and received
    Address family IPv6 Unicast: advertised and received
    Address family VPNv6 Unicast: advertised and received
    Address family IPv6 Labeled Unicast: advertised and received
  Received 3244 messages, 2 notifications, 0 in queue
  Sent 3242 messages, 2 notifications, 0 in queue
  Route refresh request: received 0, sent 2
  Minimum time between advertisement runs is 0 seconds
  Update source is lo

For address family: IPv4 Unicast  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  1 accepted prefixes
  4 announced prefixes

For address family: VPNv4 Unicast  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  AF-dependant capabilities:
    Outbound Route Filter (ORF) type (64) Prefix-list:
      Send-mode: advertised
      Receive-mode: received
    Outbound Route Filter (ORF) type (128) Prefix-list:
      Send-mode: advertised
      Receive-mode: received
  Outbound Route Filter (ORF): sent;
  Route-Reflector Client
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  Inbound path policy configured
  Incoming update prefix filter list is *ORF1
  1 accepted prefixes
  5 announced prefixes

For address family: IPv4 Labeled-Unicast  BGP table version 6, neighbor version
6
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Route-Reflector Client
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  1 accepted prefixes
  4 announced prefixes

For address family: L2VPN VPLS  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
```

```
  9 announced prefixes

For address family: L2VPN EVPN  BGP table version 5, neighbor version 5
  Index 1, Offset 0, Mask 0x2
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  20 accepted prefixes
  Accepted AD:10 MACIP:0 MCAST:10 ESI:0 PREFIX:0
  20 announced prefixes

For address family: IPv6 Unicast  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  0 announced prefixes

For address family: VPNv6 Unicast  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  5 accepted prefixes
  5 announced prefixes

For address family: 6PE Labeled Unicast  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  10 accepted prefixes
  10 announced prefixes

Connections established 4; dropped 3
Local host: 192.168.1.4, Local port: 44897
Foreign host: 192.168.1.1, Foreign port: 179
TCP MSS: (0), Advertise TCP MSS: (1460), Send TCP MSS: (1460),  Receive TCP MSS:
 (1460)
Sock FD : (26)
Nexthop: 192.168.1.4
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:06:19, due to BGP Notification received
Notification Error Message: (Cease/Administratively Reset.)

RR1-7038#
```

Glossary

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

Key Terms/Acronym	Description
ORF	ORF stands for Outbound Route Filtering. It is a feature in routing protocols, particularly in BGP (Border Gateway Protocol), that allows a router to advertise to its neighbor routers the set of routes it can accept or reject.
LDP	LDP stands for Label Distribution Protocol. It is a signaling protocol used in MPLS

Key Terms/Acronym	Description
	(Multiprotocol Label Switching) networks to distribute and exchange labels between MPLS-enabled routers.
OSPF	OSPF stands for Open Shortest Path First. It is a routing protocol used in computer networks, particularly in large enterprise and service provider networks.

© 2026 IP Infusion Inc. Proprietary

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 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.

Notes:

- 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.

BGP Bogon Prefix Filtering 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

© 2026 IP Infusion Inc. Proprietary

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.

Figure 94. BGP Bogon Filter

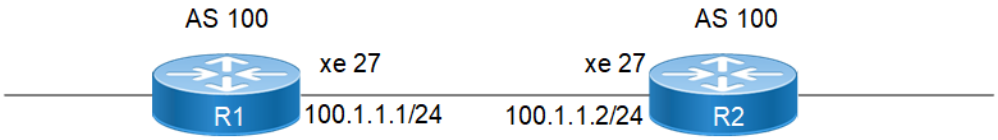

Prerequisites

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

```
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

```
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
```

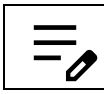

Note: Before configuration, meet all Prerequisites (page 1149).

Configure Invalid IP Addresses on R1 and R2

1. 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
```

2. 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.

Notes:

- 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.

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.

© 2026 IP Infusion Inc. Proprietary

BGP Large Communities

Overview

BGP routing policies control route distribution and behavior across networks. Communities, optional transitive BGP attributes associate operational metadata with routes, like origin location or specific routing actions. However, standard and extended BGP community attributes fall short in environments using 4-byte Autonomous System Numbers (ASNs).

To solve this issue, BGP introduced the Large Community (LC) attribute. It consists of an unordered set of one or more twelve-octet values, including: A 4-byte Global Administrator field (typically the ASN) and two 4-byte operator-defined fields.

These attributes enable operators to encode information pertinent to their routing policies and operational needs. Operators must not transmit duplicate values, and receiving BGP speakers must silently discard any redundant entries.

Feature Characteristics

OcNOS enhanced support for BGP Large Communities by:

- Allowing route-maps to match and manipulate multiple large-community values.
- Supporting configuration adjustments for additive operations (appending values) and deleting specific values.
- Introducing ip large-community-list (page 1173) to organize and reuse large-community match criteria.

Note: For details on upgrading and downgrading versions, refer to the *OcNOS Migration document*.

Benefits

Scalable Policy Management: Named community-lists simplify reuse and maintenance across multiple route-maps.

Enhanced Flexibility: Additive and deletion capabilities allow policy controls without impacting other community values.

Future-Proofing: Prepares networks for advanced policy frameworks as multi-ASN and operator-specific routing needs evolve.

BGP Large Communities Configuration

This section illustrates the use of BGP Large Communities for scalable and consistent route policy enforcement in a service provider or multi-tenant environment. It highlights how route tagging at the source (Router A), centralized policy processing at a Route Reflector (Router B), and controlled propagation to downstream routers (Router C) enable flexible traffic engineering, service differentiation, or route filtering, without the need for policy replication on every node.

© 2026 IP Infusion Inc. Proprietary

Topology

The figure below consists of three BGP routers in a route-reflector setup. Router B acts as the Route Reflector (RR), while routers A and C are its RR clients. There are no direct iBGP sessions between A and C; all BGP routing information exchange flows through the RR.

Each client establishes iBGP sessions with B for the following address families:

- IPv4 Unicast
- IPv6 Unicast
- VPNv4 Unicast

In this setup, Router A originates prefixes with large communities. Router B uses a policy engine to match, strip, and overwrite these communities using a route-map. It then advertises the modified routes to Router C. Router C applies a large-community filter to accept only those routes that match the expected community values. This demonstrates policy-based route control using BGP Large Communities in an RR-based MPLS network.

Figure 95. BGP Large Community-Based Routing Policy Flow

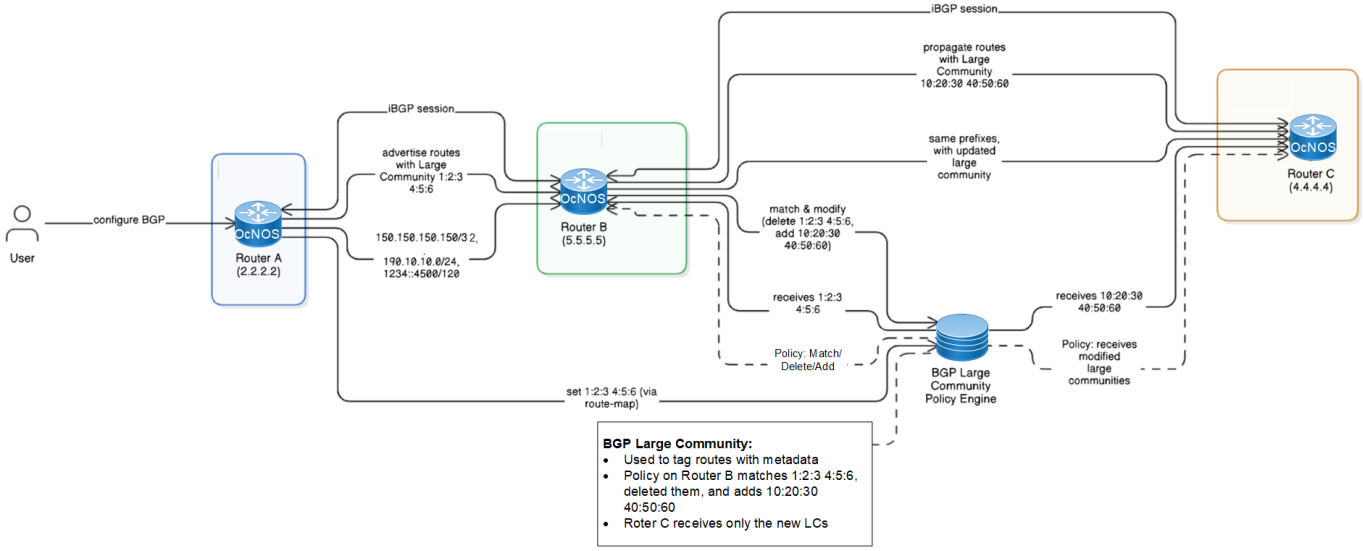

Prerequisites

Routing Configurations On Routers A, B, and C

- Set a loopback-based stable router ID for BGP, OSPF, and LDP.
- Define VRFs and import or export route targets.
- Configure LDP

Router A

```
router-id 2.2.2.2
!
```

```
ip vrf A
 rd 4200000000:2
 route-target both 1:1
 route-target both 2:2
 route-target both 3:3
 route-target both 4:4
 route-target both 5:5
 route-target both 6:6
 route-target both 7:7
 route-target both 8:8
 route-target both 9:9
!
router ldp
 transport-address ipv4 2.2.2.2
!
```

Router B

```
router-id 5.5.5.5
!
ip vrf A
 rd 4200000000:5
 route-target both 1:1
 route-target both 2:2
 route-target both 3:3
 route-target both 4:4
 route-target both 5:5
 route-target both 6:6
 route-target both 7:7
 route-target both 8:8
 route-target both 9:9
!
router ldp
 transport-address ipv4 5.5.5.5
!
```

Router C

```
router-id 4.4.4.4
!
ip vrf A
 rd 4200000001:4
 route-target both 1:1
 route-target both 2:2
 route-target both 3:3
 route-target both 4:4
 route-target both 5:5
 route-target both 6:6
 route-target both 7:7
 route-target both 8:8
 route-target both 9:9
!
router ldp
 transport-address ipv4 4.4.4.4
!
```

Interface Setup Configurations on Routers A, B, and C

- Configure the loopback interface for IPv4 or IPv6.
- Configure key interfaces connecting to core, CE, and PE neighbors. Ensure proper VRF bindings and IP assignments.

© 2026 IP Infusion Inc. Proprietary

Router A

```
interface eth8
 ip vrf forwarding A
 ip address 20.2.0.2/24
!
interface eth7
 ip address 10.17.0.2/24
!
interface eth6
 ip vrf forwarding B
 ip address 10.19.0.2/24
!
interface eth5
 ip address 10.14.0.2/24
!
interface eth4
 ip address 10.11.0.2/24
 ipv6 address ::10.11.0.2/120
 ipv6 address 2255::2522/120
 ipv6 router ospf area 0.0.0.0 instance-id 0
!
interface eth3
 ip vrf forwarding A
 ip address 10.3.0.2/24
!
interface eth2
 ip address 10.7.0.2/24
 ipv6 address ::10.7.0.2/120
 ipv6 address 2244::2422/120
 label-switching
 ipv6 router ospf area 0.0.0.0 instance-id 0
 enable-ldp ipv4
!
interface eth1
 ip vrf forwarding A
 ip address 10.1.0.2/24
 ipv6 address ::10.1.0.2/120
 ipv6 address 1122::1222/120
!
interface eth0
 ip address 10.16.42.207/24
!
interface lo
 ip address 127.0.0.1/8
 ip address 2.2.2.2/32 secondary
 ipv6 address ::1/128
 ipv6 address 2222::2222/128
 ipv6 router ospf area 0.0.0.0 instance-id 0
!
 exit
!
```

Router B

```
interface eth6
 ip address 10.10.0.5/24
 label-switching
 ipv6 router ospf area 0.0.0.0 instance-id 0
 enable-ldp ipv4
!
interface eth5
!
interface eth4
 ip address 10.11.0.5/24
 ipv6 address ::10.11.0.0/24
```

© 2026 IP Infusion Inc. Proprietary

```
 ipv6 address 2525::2555/120
 ipv6 router ospf area 0.0.0.0 instance-id 0
!
interface eth3
 ip vrf forwarding A
 ip address 20.5.0.5/24
!
interface eth2
!
interface eth1
!
interface eth0
 ip address 10.16.42.210/24
!
interface lo
 ip address 127.0.0.1/8
 ip address 5.5.5.5/32 secondary
 ipv6 address ::1/128
 ipv6 address 5555::5555/128
 ipv6 router ospf area 0.0.0.0 instance-id 0
!
 exit
!
```

Router C

```
interface eth6
 ip address 10.10.0.4/24
 ipv6 address ::10.10.0.4/120
 ipv6 address 4545::4544/120
 label-switching
 ipv6 router ospf area 0.0.0.0 instance-id 0
 enable-ldp ipv4
!
interface eth5
 ip address 10.5.0.4/24
 label-switching
 enable-ldp ipv4
!
interface eth4
 ip address 10.9.0.4/24
 label-switching
 enable-ldp ipv4
!
interface eth3
 ip address 10.8.0.4/24
 label-switching
 enable-ldp ipv4
!
interface eth2
 ip address 10.7.0.4/24
 ipv6 address ::10.7.0.4/120
 ipv6 address 2244::2444/120
 label-switching
 ipv6 router ospf area 0.0.0.0 instance-id 0
 enable-ldp ipv4
!
interface eth1
 ip vrf forwarding A
 ip address 20.4.0.4/24
!
interface eth0
 ip address 10.16.42.209/24
!
interface lo
 ip address 127.0.0.1/8
 ip address 4.4.4.4/32 secondary
 ipv6 address ::1/128
```

© 2026 IP Infusion Inc. Proprietary

```
 ipv6 address 4444::4444/128
 ipv6 router ospf area 0.0.0.0 instance-id 0
!
 exit
!
```

Configure OSPF on Routers A, B, and C

Router A

```
router ospf 1
 network 2.2.2.2/32 area 0.0.0.0
 network 10.7.0.0/24 area 0.0.0.0
 network 10.11.0.0/24 area 0.0.0.0
 network 10.14.0.0/24 area 0.0.0.0
 network 10.17.0.0/24 area 0.0.0.0
 network 10.19.0.0/24 area 0.0.0.0
!
router ipv6 ospf
!
```

Router B

```
router ospf 1
 network 5.5.5.5/32 area 0.0.0.0
 network 10.10.0.0/24 area 0.0.0.0
 network 10.11.0.0/24 area 0.0.0.0
!
router ipv6 ospf
!
```

Router C

```
router ospf 1
 network 4.4.4.4/32 area 0.0.0.0
 network 10.5.0.0/24 area 0.0.0.0
 network 10.7.0.0/24 area 0.0.0.0
 network 10.8.0.0/24 area 0.0.0.0
 network 10.9.0.0/24 area 0.0.0.0
 network 10.10.0.0/24 area 0.0.0.0
!
router ipv6 ospf
!
```

Configure Static Routes on Routers A and C

Ensure fallback reachability.

Router A

```
ip route 150.150.150.150/32 10.14.0.150
!
ipv6 route 1234::4500/120 2244::2423
!
```

Router C

```
!
ip route vrf A 60.10.10.0/24 20.4.0.60
!
```

© 2026 IP Infusion Inc. Proprietary

BGP Large Communities Routers Configurations

Note: Before configuration, meet all Prerequisites (page 1154).

Router A - Large Community Configuration

Tag outbound BGP routes with two large community values (`1:2:3 and 4:5:6`) to influence downstream routing behavior.

Define the Route Map to Set Communities

Router A adds large community values `1:2:3 and 4:5:6` to outbound routes.

```
OcNOS(config)#route-map set-lc-tag-out permit 10
OcNOS(config-route-map)#set large-community 1:2:3 4:5:6
OcNOS(config-route-map)#commit
OcNOS(config-route-map)#exit
```

Apply the Route Map in BGP

To all relevant address-families (IPv4, VPNv4, IPv6), apply the route map outbound to peer B (`5.5.5.5 and 5555::5555`)

```
router bgp 4200000000
 neighbor 5.5.5.5 remote-as 4200000000
 neighbor 5.5.5.5 update-source lo
 neighbor 5555::5555 remote-as 4200000000
 neighbor 5555::5555 update-source lo
 !
 address-family ipv4 unicast
 redistribute static
 neighbor 5.5.5.5 activate
 neighbor 5.5.5.5 route-map set-lc-tag-out out
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 5.5.5.5 activate
 neighbor 5.5.5.5 route-map set-lc-tag-out out
 exit-address-family
 !
 address-family ipv6 unicast
 redistribute static
 neighbor 5555::5555 activate
 neighbor 5555::5555 route-map set-lc-tag-out out
 exit-address-family
 !
 exit
!
```

Router B - Large Community Match and Rewrite Configuration

Match received BGP routes containing large community `1:2:3 and 4:5:6`, strip them, and apply new values `40:50:60 and 10:20:30`.

Define a Large Community List

```
OcNOS(config)#ip large-community-list 1 permit 1:2:3 4:5:6
OcNOS(config)#commit
```

© 2026 IP Infusion Inc. Proprietary

Create a Route Map with Match and Rewrite

Router B detects and removes original communities, then re-tags the route with updated communities.

```
OcNOS(config)#route-map set-lc-tag-out permit 10
OcNOS(config-route-map)#match large-community 1
OcNOS(config-route-map)#set large-comm-list 1 delete
OcNOS(config-route-map)#set large-community 40:50:60 10:20:30 additive
OcNOS(config-route-map)#commit
OcNOS(config-route-map)#exit
```

Apply the Route Map in BGP

Apply route-map outbound only on peer C (`4.4.4.4 and 4444::4444`).

```
router bgp 4200000000
 neighbor 2.2.2.2 remote-as 4200000000
 neighbor 2.2.2.2 update-source lo
 neighbor 4.4.4.4 remote-as 4200000000
 neighbor 4.4.4.4 update-source lo
 neighbor 2222::2222 remote-as 4200000000
 neighbor 2222::2222 update-source lo
 neighbor 4444::4444 remote-as 4200000000
 neighbor 4444::4444 update-source lo
 !
 address-family ipv4 unicast
 neighbor 2.2.2.2 activate
 neighbor 2.2.2.2 route-reflector-client
 neighbor 4.4.4.4 activate
 neighbor 4.4.4.4 route-reflector-client
 neighbor 4.4.4.4 route-map set-lc-tag-out out
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 2.2.2.2 activate
 neighbor 2.2.2.2 route-reflector-client
 neighbor 4.4.4.4 activate
 neighbor 4.4.4.4 route-reflector-client
 neighbor 4.4.4.4 route-map set-lc-tag-out out
 exit-address-family
 !
 address-family ipv6 unicast
 neighbor 2222::2222 activate
 neighbor 2222::2222 route-reflector-client
 neighbor 4444::4444 activate
 neighbor 4444::4444 route-reflector-client
 neighbor 4444::4444 route-map set-lc-tag-out out
 exit-address-family
 !
 exit
!
```

Router C - Passive Role in Large Community Policy

Participate in BGP flow and MPLS forwarding. No large community changes or filters. Acts as a transit or receiving node.

- No `ip large-community-list` or `route-map` configuration.
- Standard BGP peerings with B (`5.5.5.5 and 5555::5555`).
- Router C accepts routes tagged with modified large community from B and may forward based on policy.

```
router bgp 4200000000
 neighbor 5.5.5.5 remote-as 4200000000
```

© 2026 IP Infusion Inc. Proprietary

```
 neighbor 5.5.5.5 update-source lo
 neighbor 5555::5555 remote-as 4200000000
 neighbor 5555::5555 update-source lo
 !
 address-family ipv4 unicast
 neighbor 5.5.5.5 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 5.5.5.5 activate
 exit-address-family
 !
 address-family ipv6 unicast
 neighbor 5555::5555 activate
 exit-address-family
 !
 address-family ipv4 vrf A
 exit-address-family
 !
 exit
!
ip route vrf A 60.10.10.0/24 20.4.0.60
!
```

BGP Large Communities Validation

Verify that the routes are being advertised correctly to Router B and are carrying the intended community tags (in this case, Large Communities 1:2:3 and 4:5:6).

The `show ip bgp` command confirms that the IPv4 unicast route (e.g. `150.150.150.150/32`) exists in the local BGP table. The BGP session to B (`neighbor 5.5.5.5`) is active and routes are sent.

```
Router A#show ip bgp
BGP table version is 1, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
              l - labeled, S Stale, x-EVPN
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

    Network          Next Hop            Metric    LocPrf   Weight Path   Ext-Color
*>   150.150.150.150/32
                      10.14.0.150          0        100       32768  ?        -

Total number of prefixes 1

Router B#show ip bgp
BGP table version is 3, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
              l - labeled, S Stale, x-EVPN
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

    Network          Next Hop            Metric    LocPrf   Weight Path   Ext-Color
*>i  150.150.150.150/32
                      10.14.0.150          0        100       0    ?        -

Total number of prefixes 1
```

The `show ip bgp 150.150.150.150` command confirms that the route is locally originated and advertised to B. It is useful for verifying if the prefix is a candidate for applying the Large Community policy.

```
Router A#show ip bgp 150.150.150.150
BGP routing table entry for 150.150.150.150/32
```

```
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  5.5.5.5
  Local
  Path Selection reason: Nothing left to compare
   Nexthop:10.14.0.150 from 0.0.0.0 (Router ID:2.2.2.2)
      Origin incomplete,metric 0, localpref 100, weight 32768      valid, sourced, best, source-safi:
1
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Last update: Tue Jun  3 12:18:26 2025, 00:41:49 ago

Router B#show ip bgp 150.150.150.150
BGP routing table entry for 150.150.150.150/32
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  4.4.4.4
  Local, (Received from a RR-client)
  Path Selection reason: Nothing left to compare
   Nexthop:10.14.0.150 (IGP metric 2) from 2.2.2.2 (Remote Id:2.2.2.2) Peer nexthop: 2.2.2.2
      Origin incomplete, metric 0, localpref 100      valid, internal, best, source-safi: 1
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Large Community: 1:2:3 4:5:6
      Last update: Tue Jun  3 12:55:57 2025, 00:01:18 ago

Router C#show ip bgp 150.150.150.150
BGP routing table entry for 150.150.150.150/32
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  Local
  Path Selection reason: Nothing left to compare
   Nexthop:10.14.0.150 (IGP metric 3) from 5.5.5.5 (Originator Id:2.2.2.2) (Remote Id:5.5.5.5) Peer
nexthop: 5.5.5.5
      Origin incomplete, metric 0, localpref 100      valid, internal, best, source-safi: 1
      Originator: 2.2.2.2, Cluster list: 5.5.5.5
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Large Community: 10:20:30 40:50:60
      Last update: Tue Jun  3 12:55:57 2025, 00:01:35 ago
```

The command `show ip bgp vpnv4 vrf A` confirms that the routes are originated and tagged with route targets, and the redistribution from VRF A to BGP VPNv4 is occurring.

```
Router A#show ip bgp vpnv4 vrf A
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
Route Distinguisher: 4200000000:2 (Default for VRF A)
*> l 190.10.10.0/24   10.1.0.1             0        100       0   4200000002 ?        -
 Announced routes count = 1
 Accepted routes count = 0

Router B#show ip bgp vpnv4 all
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
Route Distinguisher: -94967296:2
*>i  190.10.10.0/24   2.2.2.2              0        100       0   4200000002 ?        -
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: -94967296:5 (Default for VRF A)
*>i  190.10.10.0/24   2.2.2.2              0        100       0   4200000002 ?        -
```

```
 Announced routes count = 0
 Accepted routes count = 1
```

The command `show ip bgp vpnv4 vrf A 190.10.10.0` confirms that Router A is advertising VPNv4 routes (e.g. 190.10.10.0/24) to Router B and VRF-specific community values like RTs.

```
Router A#show ip bgp vpnv4 vrf A 190.10.10.0
Route Distinguisher: 4200000000:2 (Default for VRF A) Routing Entry for prefix: 190.10.10.0/24
  Advertised to non peer-group peers:
  5.5.5.5
  AS path:{4200000002}
  Path Selection reason: Nothing left to compare
   Nexthop:10.1.0.1 (IGP metric 0) from 10.1.0.1 (Remote Id:10.16.42.206) Peer nexthop: 10.1.0.1
      Origin incomplete, metric 0, localpref 100, Out-label 0, In-label 24320, refcnt: 1
      valid, external, best, source-safi: 1
      Extended Community: RT:1:1 2:2 3:3 4:4 5:5 6:6 7:7 8:8 9:9
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Last update: Mon Jun  2 07:53:33 2025, 1d05h06m ago

Router B#show ip bgp vpnv4 all 190.10.10.0
Route Distinguisher: -94967296:2 Routing Entry for prefix: 190.10.10.0/24
  Advertised to non peer-group peers:
  4.4.4.4
  AS path:{4200000002}, (Received from a RR-client)
  Path Selection reason: Nothing left to compare
   Nexthop:2.2.2.2 (IGP metric 2) from 2.2.2.2 (Remote Id:2.2.2.2) Peer nexthop: 2.2.2.2
      Origin incomplete, metric 0, localpref 100, Out-label 24320      valid, internal, best, source-
safi: 128
      Extended Community: RT:1:1 2:2 3:3 4:4 5:5 6:6 7:7 8:8 9:9
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Large Community: 1:2:3 4:5:6
      Last update: Tue Jun  3 12:55:54 2025, 00:02:02 ago

Route Distinguisher: -94967296:5 (Default for VRF A) Routing Entry for prefix: 190.10.10.0/24
  Not advertised to any peer
  AS path:{4200000002}, (Received from a RR-client)
  Path Selection reason: Nothing left to compare
   Nexthop:2.2.2.2 (IGP metric 2) from 2.2.2.2 (Remote Id:2.2.2.2) Peer nexthop: 2.2.2.2
      Origin incomplete, metric 0, localpref 100, Out-label 24320      valid, internal, best, source-
safi: 128
      Duplicated: (source VRF-ID: 0, source VRF: DEFAULT, VRF-External, imported)
      Extended Community: RT:1:1 2:2 3:3 4:4 5:5 6:6 7:7 8:8 9:9
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Large Community: 1:2:3 4:5:6
      Last update: Tue Jun  3 12:55:54 2025, 00:02:02 ago

 Router C#show ip bgp vpnv4 vrf A 190.10.10.0
Route Distinguisher: -94967295:4 (Default for VRF A) Routing Entry for prefix: 190.10.10.0/24
  Not advertised to any peer
  AS path:{4200000002}
  Path Selection reason: Nothing left to compare
   Nexthop:2.2.2.2 (IGP metric 3) from 5.5.5.5 (Originator Id:2.2.2.2) (Remote Id:5.5.5.5) Peer
nexthop: 5.5.5.5
      Origin incomplete, metric 0, localpref 100, Out-label 24320      valid, internal, best, source-
safi: 128
      Duplicated: (source VRF-ID: 0, source VRF: DEFAULT, VRF-External, imported)
      Extended Community: RT:1:1 2:2 3:3 4:4 5:5 6:6 7:7 8:8 9:9
      Originator: 2.2.2.2, Cluster list: 5.5.5.5
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Large Community: 10:20:30 40:50:60
      Last update: Tue Jun  3 12:55:57 2025, 00:02:19 ago
```

The command `show bgp ipv6` confirms that IPv6 prefixes (e.g., 1234::4500/120) are present, valid, and advertised to B. It verifies that the IPv6 session to 5555::5555 is working and that the route-map is applied.

© 2026 IP Infusion Inc. Proprietary

```
Router A#show bgp ipv6
BGP table version is 1, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
              l - labeled, S Stale, x-EVPN
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

    Network          Next Hop                    Metric    LocPrf     Weight Path
*>   1234::4500/120   2244::2423                     0        100       32768  ?        -

Total number of prefixes 1

Router B#show bgp ipv6
BGP table version is 3, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
              l - labeled, S Stale, x-EVPN
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

    Network          Next Hop                    Metric    LocPrf     Weight Path
*>i  1234::4500/120   2244::2423                     0        100       0    ?        -

Total number of prefixes 1

Router A#show bgp ipv6 1234::4500
BGP routing table entry for 1234::4500/120
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  5555::5555
  Local
  Path Selection reason: Nothing left to compare
   Nexthop:2244::2423 from :: (Router ID:2.2.2.2)
      Origin incomplete,metric 0, localpref 100, weight 32768      valid, sourced, best, source-safi:
1
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Last update: Tue Jun  3 12:06:27 2025, 00:54:26 ago

Router B#show bgp ipv6 1234::4500
BGP routing table entry for 1234::4500/120
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  4444::4444
  Local, (Received from a RR-client)
  Path Selection reason: Nothing left to compare
   Nexthop:2244::2423 (IGP metric 2) from 2222::2222 (Remote Id:2.2.2.2) Peer nexthop: 2222::2222
      Origin incomplete, metric 0, localpref 100      valid, internal, best, source-safi: 1
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Large Community: 1:2:3 4:5:6
      Last update: Tue Jun  3 12:55:58 2025, 00:02:39 ago

Router C#show bgp ipv6 1234::4500
BGP routing table entry for 1234::4500/120
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  Local
  Path Selection reason: Nothing left to compare
   Nexthop:2244::2423 (IGP metric 0) from 5555::5555 (Originator Id:2.2.2.2) (Remote Id:5.5.5.5) Peer
nexthop: 5555::5555
      Origin incomplete, metric 0, localpref 100      valid, internal, best, source-safi: 1
      Originator: 2.2.2.2, Cluster list: 5.5.5.5
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Large Community: 10:20:30 40:50:60
      Last update: Tue Jun  3 12:56:01 2025, 00:02:45 ago
```

© 2026 IP Infusion Inc. Proprietary

Running Configurations

Router A Running Configuration

```
Router A#show running-config
!
! Software version: XP-6.6.1 06/02/2025 07:37:20
!
! Last configuration change at 12:19:38 UTC Tue Jun 03 2025 by root
!
feature netconf-ssh
feature netconf-tls
!
service password-encryption
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
bgp extended-asn-cap
!
qos enable
!
hostname Router A
errdisable cause stp-bpdu-guard
feature telnet
feature ssh
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp
ntp enable
!
router-id 2.2.2.2
!
ip vrf A
 rd 4200000000:2
 route-target both 1:1
 route-target both 2:2
 route-target both 3:3
 route-target both 4:4
 route-target both 5:5
 route-target both 6:6
 route-target both 7:7
 route-target both 8:8
 route-target both 9:9
!
router ldp
 transport-address ipv4 2.2.2.2
!
route-map set-lc-tag-out permit 10
 set large-community 1:2:3 4:5:6
!
interface eth8
 ip vrf forwarding A
 ip address 20.2.0.2/24
!
interface eth7
 ip address 10.17.0.2/24
!
interface eth6
 ip vrf forwarding B
 ip address 10.19.0.2/24
!
interface eth5
 ip address 10.14.0.2/24
!
interface eth4
```

© 2026 IP Infusion Inc. Proprietary

```
 ip address 10.11.0.2/24
 ipv6 address ::10.11.0.2/120
 ipv6 address 2255::2522/120
 ipv6 router ospf area 0.0.0.0 instance-id 0
!
interface eth3
 ip vrf forwarding A
 ip address 10.3.0.2/24
!
interface eth2
 ip address 10.7.0.2/24
 ipv6 address ::10.7.0.2/120
 ipv6 address 2244::2422/120
 label-switching
 ipv6 router ospf area 0.0.0.0 instance-id 0
 enable-ldp ipv4
!
interface eth1
 ip vrf forwarding A
 ip address 10.1.0.2/24
 ipv6 address ::10.1.0.2/120
 ipv6 address 1122::1222/120
!
interface eth0
 ip address 10.16.42.207/24
!
interface lo
 ip address 127.0.0.1/8
 ip address 2.2.2.2/32 secondary
 ipv6 address ::1/128
 ipv6 address 2222::2222/128
 ipv6 router ospf area 0.0.0.0 instance-id 0
!
 exit
!
router ospf 1
 network 2.2.2.2/32 area 0.0.0.0
 network 10.7.0.0/24 area 0.0.0.0
 network 10.11.0.0/24 area 0.0.0.0
 network 10.14.0.0/24 area 0.0.0.0
 network 10.17.0.0/24 area 0.0.0.0
 network 10.19.0.0/24 area 0.0.0.0
!
router ipv6 ospf
!
bgp nexthop-trigger enable
bgp nexthop-trigger delay 1
!
router bgp 4200000000
 neighbor 5.5.5.5 remote-as 4200000000
 neighbor 5.5.5.5 update-source lo
 neighbor 5555::5555 remote-as 4200000000
 neighbor 5555::5555 update-source lo
 !
 address-family ipv4 unicast
 redistribute static
 neighbor 5.5.5.5 activate
 neighbor 5.5.5.5 route-map set-lc-tag-out out
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 5.5.5.5 activate
 neighbor 5.5.5.5 route-map set-lc-tag-out out
 exit-address-family
 !
 address-family ipv6 unicast
 redistribute static
 neighbor 5555::5555 activate
```

```
 neighbor 5555::5555 route-map set-lc-tag-out out
 exit-address-family
 !
 address-family ipv4 vrf A
 neighbor 10.1.0.1 remote-as 4200000002
 neighbor 10.1.0.1 activate
 exit-address-family
 !
 exit
!
ip route 150.150.150.150/32 10.14.0.150
!
ipv6 route 1234::4500/120 2244::2423
!
line vty 0
 exec-timeout 0 0
!
!
end
```

Router B Running Configuration

```
Router B#show running-config
!
! Software version: DEMO_VM-OcNOS-SP-MPLS-x86-6.6.1.113-Alpha 05/26/2025 20:27:01
!
! Last configuration change at 12:56:46 UTC Tue Jun 03 2025 by root
!
feature netconf-ssh
feature netconf-tls
!
service password-encryption
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
bgp extended-asn-cap
!
qos enable
!
hostname Router B
errdisable cause stp-bpdu-guard
feature telnet
feature ssh
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp
ntp enable
!
router-id 5.5.5.5
!
ip vrf A
 rd 4200000000:5
 route-target both 1:1
 route-target both 2:2
 route-target both 3:3
 route-target both 4:4
 route-target both 5:5
 route-target both 6:6
 route-target both 7:7
 route-target both 8:8
 route-target both 9:9
!
ip vrf management
!
router ldp
```

```
 transport-address ipv4 5.5.5.5
!
route-map set-lc-tag-out permit 10
 match large-community 1
 set large-comm-list 1 delete
 set large-community 40:50:60 10:20:30 additive
!
interface eth8
!
interface eth7
!
interface eth6
 ip address 10.10.0.5/24
 label-switching
 ipv6 router ospf area 0.0.0.0 instance-id 0
 enable-ldp ipv4
!
interface eth5
!
interface eth4
 ip address 10.11.0.5/24
 ipv6 address ::10.11.0.0/24
 ipv6 address 2525::2555/120
 ipv6 router ospf area 0.0.0.0 instance-id 0
!
interface eth3
 ip vrf forwarding A
 ip address 20.5.0.5/24
!
interface eth2
!
interface eth1
!
interface eth0
 ip address 10.16.42.210/24
!
interface lo
 ip address 127.0.0.1/8
 ip address 5.5.5.5/32 secondary
 ipv6 address ::1/128
 ipv6 address 5555::5555/128
 ipv6 router ospf area 0.0.0.0 instance-id 0
!
 exit
!
router ospf 1
 network 5.5.5.5/32 area 0.0.0.0
 network 10.10.0.0/24 area 0.0.0.0
 network 10.11.0.0/24 area 0.0.0.0
!
router ipv6 ospf
!
router bgp 4200000000
 neighbor 2.2.2.2 remote-as 4200000000
 neighbor 2.2.2.2 update-source lo
 neighbor 4.4.4.4 remote-as 4200000000
 neighbor 4.4.4.4 update-source lo
 neighbor 2222::2222 remote-as 4200000000
 neighbor 2222::2222 update-source lo
 neighbor 4444::4444 remote-as 4200000000
 neighbor 4444::4444 update-source lo
 !
 address-family ipv4 unicast
 neighbor 2.2.2.2 activate
 neighbor 2.2.2.2 route-reflector-client
 neighbor 4.4.4.4 activate
 neighbor 4.4.4.4 route-reflector-client
 neighbor 4.4.4.4 route-map set-lc-tag-out out
```

© 2026 IP Infusion Inc. Proprietary

```
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 2.2.2.2 activate
 neighbor 2.2.2.2 route-reflector-client
 neighbor 4.4.4.4 activate
 neighbor 4.4.4.4 route-reflector-client
 neighbor 4.4.4.4 route-map set-lc-tag-out out
 exit-address-family
 !
 address-family ipv6 unicast
 neighbor 2222::2222 activate
 neighbor 2222::2222 route-reflector-client
 neighbor 4444::4444 activate
 neighbor 4444::4444 route-reflector-client
 neighbor 4444::4444 route-map set-lc-tag-out out
 exit-address-family
 !
 exit
!
ip large-community-list 1 permit 1:2:3 4:5:6
!
!
end
```

Router C Running Configuration

```
Router C#show running-config
!
! Software version: DEMO_VM-OcNOS-SP-MPLS-x86-6.6.1.120-Alpha 06/01/2025 20:26:21
!
! Last configuration change at 12:21:29 UTC Tue Jun 03 2025 by root
!
feature netconf-ssh
feature netconf-tls
!
service password-encryption
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
bgp extended-asn-cap
!
qos enable
!
hostname Router C
errdisable cause stp-bpdu-guard
feature telnet
feature ssh
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp
ntp enable
!
router-id 4.4.4.4
!
ip vrf A
 rd 4200000001:4
 route-target both 1:1
 route-target both 2:2
 route-target both 3:3
 route-target both 4:4
 route-target both 5:5
 route-target both 6:6
 route-target both 7:7
```

```
 route-target both 8:8
 route-target both 9:9
!
ip vrf management
!
router ldp
 transport-address ipv4 4.4.4.4
!
interface eth8
!
interface eth7
!
interface eth6
 ip address 10.10.0.4/24
 ipv6 address ::10.10.0.4/120
 ipv6 address 4545::4544/120
 label-switching
 ipv6 router ospf area 0.0.0.0 instance-id 0
 enable-ldp ipv4
!
interface eth5
 ip address 10.5.0.4/24
 label-switching
 enable-ldp ipv4
!
interface eth4
 ip address 10.9.0.4/24
 label-switching
 enable-ldp ipv4
!
interface eth3
 ip address 10.8.0.4/24
 label-switching
 enable-ldp ipv4
!
interface eth2
 ip address 10.7.0.4/24
 ipv6 address ::10.7.0.4/120
 ipv6 address 2244::2444/120
 label-switching
 ipv6 router ospf area 0.0.0.0 instance-id 0
 enable-ldp ipv4
!
interface eth1
 ip vrf forwarding A
 ip address 20.4.0.4/24
!
interface eth0
 ip address 10.16.42.209/24
!
interface lo
 ip address 127.0.0.1/8
 ip address 4.4.4.4/32 secondary
 ipv6 address ::1/128
 ipv6 address 4444::4444/128
 ipv6 router ospf area 0.0.0.0 instance-id 0
!
 exit
!
router ospf 1
 network 4.4.4.4/32 area 0.0.0.0
 network 10.5.0.0/24 area 0.0.0.0
 network 10.7.0.0/24 area 0.0.0.0
 network 10.8.0.0/24 area 0.0.0.0
 network 10.9.0.0/24 area 0.0.0.0
 network 10.10.0.0/24 area 0.0.0.0
!
router ipv6 ospf
```

© 2026 IP Infusion Inc. Proprietary

```
!
bgp nexthop-trigger enable
!
router bgp 4200000000
 neighbor 5.5.5.5 remote-as 4200000000
 neighbor 5.5.5.5 update-source lo
 neighbor 5555::5555 remote-as 4200000000
 neighbor 5555::5555 update-source lo
 !
 address-family ipv4 unicast
 neighbor 5.5.5.5 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 5.5.5.5 activate
 exit-address-family
 !
 address-family ipv6 unicast
 neighbor 5555::5555 activate
 exit-address-family
 !
 address-family ipv4 vrf A
 exit-address-family
 !
 exit
!
ip route vrf A 60.10.10.0/24 20.4.0.60
!
line console 0
 exec-timeout 0 0
line vty 0
 exec-timeout 0 0
!
!
end
```

Implementation Example

Use Case: A transit provider assigns unique BGP Large Community values to each customer. These values are utilized in route maps to influence outbound policy decisions, such as local preference or AS-path modification.

Define large-community values associated with Customer A:

```
ip large-community-list standard cust_A permit 65001:100:1
ip large-community-list standard cust_A permit 65001:100:2
```

Apply policy using a route-map:

```
route-map export-cust_A permit 10
 match large-community cust_A
 set local-preference 150
```

This configuration applies a higher local preference to prefixes received from Customer A. By using a named large-community list, the setup remains scalable-community values can be updated in one place without editing the route map directly. This approach also simplifies policy management across multiple route maps or peers.

BGP Large Communities Commands

The following are the BGP LC commands.

© 2026 IP Infusion Inc. Proprietary

set large-community

- The existing syntax now includes the new `additive` parameter that allows users to append large community values to a route.
- Users can configure multiple large communities in a route map. The character limit for each community has increased from 32 to 255 characters.

For more details, refer to the set large-community (page 378) command in the section in the , OcNOS Layer 3 document.

© 2026 IP Infusion Inc. Proprietary

ip large-community-list

Use this command to define a set of BGP Large Community values for use in policy configuration (e.g., route-maps). It functions similarly to standard or extended community-lists and allows reuse of defined large-community patterns.

Use `no` parameter of this command to remove the configured large community values.

Command Syntax

```
ip large-community-list {<1-99> | standard WORD | <100-500> | expanded WORD} {permit | deny}
(.AA:NN:MM)
no ip large-community-list {<1-99> | standard WORD | <100-500> | expanded WORD} {permit | deny}
(.AA:NN:MM)
```

Parameters

<1-99>
Specifies a standard large-community-list using a numerical ID. Standard lists match exact values.

standard WORD
Specifies a named standard list that matches exact large-community values.

<100-500>
Specifies an expanded large-community-list using a numerical ID. Expanded lists support regex-style pattern matching.

expanded WORD
Specifies a named expanded list that uses pattern matching for flexible matching criteria.

permit
Permits the specified large-community value(s) to match in a policy.

deny
Denies the specified large-community value(s) from matching.

.AA:NN:MM
A 12-byte large community in the format ASN:Value1:Value2. E.g., 65001:100:1. Large Communities use the format Global Administrator:LocalData1:LocalData2, where each field is a 32-bit integer.

Default

None

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.6.1.

Example

The following example demonstrates how to configure different types of BGP Large Community lists in OcNOS and verify them using the `show` command.

Creates a standard large-community-list with numeric ID `1` and permits the community `4:5:6`.

```
OcNOS#configure terminal
OcNOS(config)#ip large-community-list 1 permit 4:5:6
OcNOS(config)#commit
```

© 2026 IP Infusion Inc. Proprietary

Creates an expanded large-community-list with numeric ID `150` and denies the community `1000:2000:3000`.

```
OcNOS(config)#ip large-community-list 150 deny 1000:2000:3000
OcNOS(config)#commit
```

Defines a named expanded list called `EXPTEST` and permits the community `5500:3000:2000`.

```
OcNOS(config)#ip large-community-list expanded EXPTEST permit 5500:3000:2000
OcNOS(config)#commit
```

Defines a named standard list called `STDTEST` and permits the community `50:50:50`.

```
OcNOS(config)#ip large-community-list standard STDTEST permit 50:50:50
OcNOS(config)#commit
```

Verify the configured lists using `show ip large-community-list` command.

```
OcNOS#show ip large-community-list
Large community standard list 1
    permit 4:5:6
Large community (expanded) list 150
    deny 1000:2000:3000
Named large community expanded list EXPTEST
    permit 5500:3000:2000
Named large community standard list STDTEST
    permit 50:50:50
```

© 2026 IP Infusion Inc. Proprietary

match large-community

Use this command to match BGP routes based on Large Community values. This allows selective policy actions (e.g., setting local preference, AS path, AIGP) for routes carrying specific BGP Large Community attributes.

Use `no` parameter of this command to remove the configured Large Community values.

Note: Only one `match large-community` is allowed per route-map sequence. A new match replaces any previous configuration in the same sequence.

Command Syntax

```
match large-community {<1-99> | <100-500> | WORD} (exact-match|)
no match large-community {<1-99> | <100-500> | WORD} (exact-match|)
```

Parameters

<1-99>
Specifies a standard large-community-list using a numerical ID. Standard lists match exact values.

<100-500>
Specifies an expanded large-community-list using a numerical ID. Expanded lists support regex-style pattern matching.

WORD
Specifies a named standard or expanded large-community list.

exact-match
(Optional) Matches routes only if the communities match exactly, with no additional entries.

Default

None

Command Mode

Route map mode

Applicability

Introduced in OcNOS version 6.6.1.

Example

The following example shows how to reference large-community in route-maps for policy enforcement.

Matches routes containing any large-community permitted by standard list `1`.

```
OcNOS#configure terminal
OcNOS(config)#route-map rmap-set-lc1
OcNOS(config-route-map)#match large-community 1
OcNOS(config-route-map)#commit
OcNOS(config-route-map)#exit
```

Matches routes based on pattern rules defined in expanded list `100`.

```
OcNOS(config)#route-map rmap-set-lc2
OcNOS(config-route-map)#match large-community 100
```

© 2026 IP Infusion Inc. Proprietary

```
OcNOS(config-route-map)#commit
OcNOS(config-route-map)#exit
```

Matches routes using named expanded list `EXPTEST`.

```
OcNOS(config)#route-map rmap-set-lc
OcNOS(config-route-map)#match large-community EXPTEST
OcNOS(config-route-map)#commit
OcNOS(config-route-map)#exit
```

Matches routes that have only the exact large-community entries defined in `STDTEST`.

```
OcNOS(config)#route-map rmap-set-lc3
OcNOS(config-route-map)#match large-community STDTEST exact-match
OcNOS(config-route-map)#commit
OcNOS(config-route-map)#exit
```

Verify the configured lists using `show` command.

```
OcNOS#show route-map
route-map rmap-set-lc, permit, sequence 10
  Match clauses:
    large-community: EXPTEST
  Set clauses:
route-map rmap-set-lc1, permit, sequence 10
  Match clauses:
    large-community: 1
  Set clauses:
route-map rmap-set-lc2, permit, sequence 10
  Match clauses:
    large-community: 100
  Set clauses:
route-map rmap-set-lc3, permit, sequence 10
  Match clauses:
    large-community: STDTEST exact-match
  Set clauses:

OcNOS#show running-config route-map
!
route-map rmap-set-lc permit 10
 match large-community EXPTEST
!
route-map rmap-set-lc1 permit 10
 match large-community 1
!
route-map rmap-set-lc2 permit 10
 match large-community 100
!
route-map rmap-set-lc3 permit 10
 match large-community STDTEST exact-match
!
```

set large-comm-list delete

Use this command within a route-map mode to delete the large-community-list values from matched routes (if present).

Use `no` parameter of this command to add the specified large-community values to matched routes.

Note: Only one set large-comm-list configuration is allowed per route-map sequence. A new one will overwrite the previous configuration.

Command Syntax

```
set large-comm-list {<1-99> | <100-500> | WORD} (delete)
no set large-comm-list {<1-99> | <100-500> | WORD} (delete)
```

Parameters

<1-99>
Specifies a standard large-community-list using a numerical ID. Standard lists match exact values.

<100-500>
Specifies an expanded large-community-list using a numerical ID. Expanded lists support regex-style pattern matching.

WORD
Specifies a named standard or expanded large-community list.

delete
Removes the matching large-community values from matched routes.

Default

None

Command Mode

Route map mode

Applicability

Introduced in OcNOS version 6.6.1.

Example

The following example shows how to remove (delete) selected communities from matched routes using the set large-comm-list command.

Deletes communities defined in list `100` from matched routes (`rmap2`).

```
OcNOS(config)#route-map rmap2
OcNOS(config-route-map)#match large-community 100
OcNOS(config-route-map)#set large-comm-list 100 delete
OcNOS(config-route-map)#commit
OcNOS(config-route-map)#exit
```

Deletes communities defined in list `10` from matched routes (`rmap`).

© 2026 IP Infusion Inc. Proprietary

```
OcNOS(config)#route-map rmap
OcNOS(config-route-map)#match large-community EXPTEST
OcNOS(config-route-map)#set large-comm-list 10 delete
OcNOS(config-route-map)#commit
OcNOS(config-route-map)#exit
```

Verify the configured lists using the `show` command.

- The command `set large-comm-list 100 delete` appears only for `rmap` and `rmap2`, indicating community deletion is configured only for those route-maps.
- When a route-map is configured with the `set large-comm-list <list> delete` command, the `show route-map` output will reflect this under the **Set clauses** field, verifying that the deletion of large-community values is active.

```
OcNOS#show route-map
route-map rmap, permit, sequence 10
  Match clauses:
    large-community: EXPTEST
  Set clauses:
    large-comm-list 10 delete
route-map rmap1, permit, sequence 10
  Match clauses:
    large-community: 1
  Set clauses:
route-map rmap2, permit, sequence 10
  Match clauses:
    large-community: 100
  Set clauses:
    large-comm-list 100 delete
route-map rmap3, permit, sequence 10
  Match clauses:
    large-community: STDTEST exact-match
  Set clauses:


OcNOS(config-route-map)#show running-config route-map
!
route-map rmap permit 10
 match large-community EXPTEST
 set large-comm-list 10 delete
!
route-map rmap1 permit 10
 match large-community 1
!
route-map rmap2 permit 10
 match large-community 100
 set large-comm-list 100 delete
!
route-map rmap3 permit 10
 match large-community STDTEST exact-match
!
```

© 2026 IP Infusion Inc. Proprietary

show ip large-community-list

Use this command to display the configured BGP IPv4 Large Community Lists. This includes both standard and expanded lists, named or numbered. The command helps verify large-community filters applied in route-maps for matching or modifying BGP route attributes.

Command Syntax

```
show ip large-community-list (<1-500> | WORD)
```

Parameters

<1-500>
Specifies numeric ID of a standard or expanded large-community-list.

WORD
Specifies a named standard or expanded large-community list.

Default

None

Command Mode

Execution mode and Privileged execution mode

Applicability

Introduced in OcNOS version 6.6.1.

Example

Displays all configured large community lists for both numbered and named, standard and expanded.

```
OcNOS#show ip large-community-list
Large community standard list 1
    permit 4:5:6
Large community (expanded) list 150
    deny 1000:2000:3000
Named large community expanded list EXPTEST
    permit 5500:3000:2000
Named large community standard list STDTEST
    permit 50:50:50
```

Displays only the expanded list with the number `150`. In the example below, it contains a single rule that denies large community `1000:2000:3000`.

```
OcNOS#show ip large-community-list 150
Large community (expanded) list 150
    deny 1000:2000:3000
```

Displays the named expanded list `EXPTEST`. In the example below, it permits the large community `5500:3000:2000`.

```
OcNOS#show ip large-community-list EXPTEST
Named large community expanded list EXPTEST
    permit 5500:3000:2000
```

© 2026 IP Infusion Inc. Proprietary

Troubleshooting

Symptom	Possible Cause	Solution
Route-map fails to match communities	The `ip large-community-list` is missing or mis-configured.	Check the configuration using `show ip large-community-list` and correct any errors.
Downgrade operation fails	The configuration includes unsupported LC features.	Revert to the previous version and remove all LC configurations before retrying the downgrade.
Routes behave unexpectedly	Duplicate LC values are being sent	Avoid transmitting redundant LC values; OcNOS automatically removes duplicates on receipt.

Glossary

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

Key Terms/Acronym	Description
Autonomous System Number (ASN)	A 2- or 4-byte unique identifier assigned to a network under a single administrative domain.
Large Community (LC)	A BGP attribute (RFC 8092) used to tag routes with metadata for policy decisions. Format: ASN:Value1:Value2.
Global Administrator (GA)	The first 4-byte field in a BGP Large Community, typically set to the ASN of the originator.
Route Reflector (RR)	A BGP router that reflects routes between RR clients to reduce iBGP session count in large networks.
RR Client	A BGP peer that receives and sends updates via the route reflector.
Internal BGP (iBGP)	BGP sessions between routers within the same AS.
Address Family Identifier (AFI) or Subsequent Address Family Identifier (SAFI)	Specifies the type of routes being carried (e.g., IPv4, VPNv4, IPv6).
Route Distinguisher (RD)	A unique identifier prepended to a prefix in MPLS VPNs to distinguish overlapping routes.
Route Target (RT)	An extended BGP community used to control the import or export of routes into or from a VRF.
BGP Attributes	Metadata in BGP updates that influence route selection or propagation (e.g., LC, RT, MED, Local Pref).
additive	A route-map action keyword that appends new LC values without removing existing ones.
delete	A route-map action keyword used to remove specific LC values from a route.

© 2026 IP Infusion Inc. Proprietary

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 on 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 (page 1371) 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 (page 1371) 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 (page 1371) CLI command has been modified in this release.

Feature Characteristics

The neighbor local-as (page 1371) 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 96. Disparate Autonomous System Number

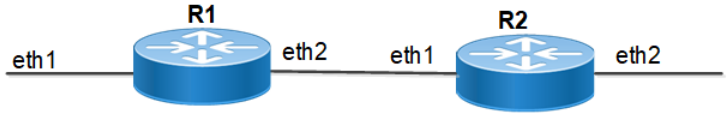

R1

`#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
```

Example

The following example show a sample configuration command.

```
#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
```

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 (page 1371)` CLI in BGP Commands (page 1244) section.

Abbreviations

Acronym	Description
ASN	Autonomous System Number
EBGP	External Border Gateway Protocol

© 2026 IP Infusion Inc. Proprietary

BGP Aggregate-Address in VRF Address-family

Overview

The aggregate-address feature for VRF address-families in BGP is designed to simplify route management by summarizing multiple specific routes into a single, broader prefix. This helps reduce the size of routing tables and minimizes the overhead for both local and remote routers. Below is an functionality explanation of the behavior and implications of this feature during configuration and un-configuration processes.

Benefits

- Route Summarization: Aggregates multiple specific routes into a single summary route, reducing routing table complexity.
- Scalability: Summarized routes optimize resource utilization, particularly in networks with multiple VRFs and large routing tables.
- Stability: Reduces route flapping as changes in individual specific routes do not impact the advertised aggregate route.

Prerequisites

- IPv4: Individual routes are advertised to BGP peers as separate prefixes.
- IPv6: Individual routes are advertised to BGP peers as separate prefixes.
- Each route is treated independently in the BGP table and shared with peers accordingly.

Configuration for Aggregate-Address in VRF Address-Family with IPv4

Note: When the VRF CE access interface flaps (via shutdown/no shutdown), connected routes become inactive, especially when OSPF is used as the PE-CE protocol. Consequently, routes learned over that interface are also inactive. To prevent the loss of these routes, configure redistribute connected under the VRF BGP instance.

Topology

This sample topology provides basic connectivity for IPv4 Aggregated address.

Figure 97. BGP Aggregate-Address in VRF Address-Family with IPv4

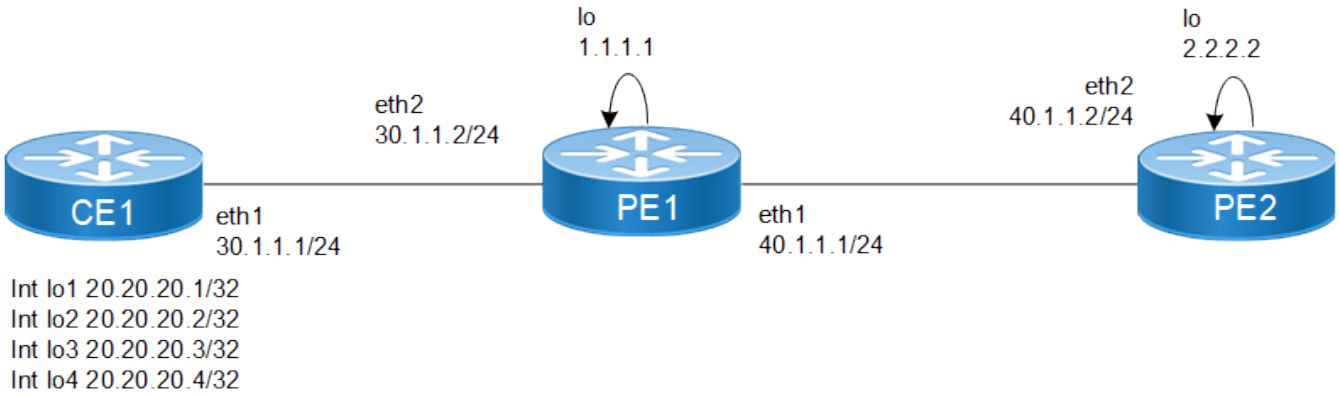

PE1

1. Configure the router ID:

```
#configure terminal
PE1 (config)#router ldp
PE1(config- router)#router-id 1.1.1.1
PE1(config- router)#exit
```

2. Create and configure VRFs:

```
PE1 (config)#ip vrf VRF1
PE1 (config-vrf)#rd 100:1
PE1 (config-vrf)#route-target both 100:1
PE1 (config-vrf)#exit
```

3. Enable LDP globally:

```
PE1(config)#router ldp
PE1(config-ldp)#exit
```

4. Configure loopback interface with IPv4 and IPv6:

```
 PE1(config-if)#interface lo
 PE1(config-if)#ip address 127.0.0.1/8
 PE1(config-if)#ip address 1.1.1.1/32 secondary
 PE1(config-if)#ipv6 address ::1/128
 PE1(config-if)#exit
```

5. Configure eth1 under VRF1 with an IP address:

```
PE1(config)#interface eth2
PE1(config-if)#ip vrf forwarding VRF1
PE1(config-if)#ip address 30.1.1.2/24
PE1(config-if)#exit
```

6. Configure eth2 for MPLS and enable LDP:

```
PE1(config)#interface eth1
PE1(config-if)#ip address 40.1.1.1/24
PE1(config-if)#ipv6 address 3000::2/64
PE1(config-if)#label-switching
PE1(config-if)#enable-ldp ipv4
PE1(config-if)#interface eth3
PE1(config-if)#interface eth4
PE1(config-if)#interface eth5
PE1(config-if)#interface eth6
PE1(config-if)#exit
```

© 2026 IP Infusion Inc. Proprietary

7. Enable OSPF and define networks:

```
PE1(config)#router ospf 100
PE1(config-router)#ospf router-id 1.1.1.1
PE1(config-router)#network 1.1.1.1/32 area 0.0.0.0
PE1(config-router)#network 40.1.1.0/24 area 0.0.0.0
PE1(config-router)#exit
```

8. Configure BGP:

```
PE1(config)#router bgp 100
PE1(config-router)#bgp router-id 1.1.1.1
PE1(config-router)#neighbor 2.2.2.2 remote-as 100
PE1(config-router)#neighbor 2.2.2.2 update-source lo
PE1(config-router)#address-family ipv4 unicast
PE1(config-router-ad)#neighbor 2.2.2.2 activate
PE1(config-router-ad)#exit-address-family
PE1(config-router)#address-family vpnv4 unicast
PE1(config-router-ad)#neighbor 2.2.2.2 activate
PE1(config-router-ad)#exit-address-family
PE1(config-router)#address-family ipv4 vrf VRF1
PE1(config-router-ad)#aggregate-address 20.20.20.0/24 as-set summary-only
PE1(config-router-ad)#neighbor 30.1.1.1 remote-as 200
PE1(config-router-ad)#neighbor 30.1.1.1 activate
PE1(config-router-ad)#neighbor 30.1.1.1 route-map map1 in
PE1(config-router-ad)#exit-address-family
```

PE2

1. Configure the router ID:

```
#configure terminal
PE2(config)#router ldp
PE2(config- router)#router-id 2.2.2.2
PE2(config- router)#exit
```

2. Create and configure VRFs:

```
PE2(config)#ip vrf VRF1
PE2(config-vrf)#rd 200:1
PE2(config-vrf)#route-target both 100:1
PE2(config-vrf)#exit
```

3. Enable LDP globally:

```
PE2(config)#router ldp
PE2(config-ldp)#exit
```

4. Configure loopback interface with IPv4 and IPv6:

```
PE2(config-if)#interface lo
PE2(config-if)#ip address 127.0.0.1/8
PE2(config-if)#ip address 2.2.2.2/32 secondary
PE2(config-if)#ipv6 address ::1/128
PE2(config-if)#exit
```

5. Configure eth1 under VRF1 with an IP address:

```
PE2(config)#interface eth1
PE2(config-if)#ip vrf forwarding VRF1
PE2(config-if)#ip address 50.1.1.2/24
PE2(config-if)#exit
```

6. Configure eth2 for MPLS and enable LDP:

```
PE2(config)#interface eth2
PE2(config-if)#ip address 40.1.1.2/24
PE2(config-if)#ipv6 address 3000::2/64
```

© 2026 IP Infusion Inc. Proprietary

```
PE2(config-if)#label-switching
PE2(config-if)#enable-ldp ipv4
PE2(config-if)#interface eth3
PE2(config-if)#interface eth4
PE2(config-if)#interface eth5
PE2(config-if)#interface eth6
PE2(config-if)#exit
```

7. Enable OSPF and define networks:

```
PE2(config)#router ospf 100
PE2(config-router)#ospf router-id 2.2.2.2
PE2(config-router)#network 2.2.2.2/32 area 0.0.0.0
PE2(config-router)#network 40.1.1.0/24 area 0.0.0.0
PE2(config-router)#exit
```

8. Configure BGP:

```
PE2(config)#router bgp 100
PE2(config-router)#bgp router-id 2.2.2.2
PE2(config-router)#neighbor 1.1.1.1 remote-as 100
PE2(config-router)#neighbor 1.1.1.1 update-source lo
PE2(config-router)#address-family ipv4 unicast
PE2(config-router-ad)#neighbor 1.1.1.1 activate
PE2(config-router-ad)#exit-address-family
PE2(config-router)#address-family vpnv4 unicast
PE2(config-router-ad)#neighbor 1.1.1.1 activate
PE2(config-router-ad)#exit-address-family
PE2(config-router)#address-family ipv4 vrf VRF1
PE2(config-router-ad)#neighbor 50.1.1.1 remote-as 200
PE2(config-router-ad)#neighbor 50.1.1.1 route-map map1 in
PE2(config-router-ad)#exit-address-family
```

CE1

1. Configure Ethernet Interface eth1 with IP Address:

```
(config)#interface eth1
CE1 (config-if)#ip address 30.1.1.1/24
CE1 (config-if)#exit
```

2. Configure Loopback Interfaces

```
(config)#interface lo1
CE1(config-if)#ip address 20.20.20.1/32
CE1(config-if)#exit
CE1(config)#interface lo2
CE1(config-if)#ip address 20.20.20.2/32
CE1(config-if)#exit
CE1(config)#interface lo3
CE1(config-if)#ip address 20.20.20.3/32
CE1(config-if)#exit
CE1(config)#interface lo4
CE1(config-if)#ip address 20.20.20.4/32
CE1(config-if)#exit
```

3. Configure BGP:

```
(config)#router bgp 300
CE1(config-router)#neighbor 30.30.30.2 remote-as 100
CE1(config-router)#address-family ipv4 unicast
CE1(config-router-ad)#neighbor 30.30.30.2 activate
CE1(config-router-ad)#redistribute connected
CE1(config-router-ad)#exit-address-family
CE1(config-router)#exit
```

© 2026 IP Infusion Inc. Proprietary

Running Configuration with IPv4

PE1

```
!
router-id 1.1.1.1
!
ip vrf management
!
ip vrf VRF1
 rd 100:1
 route-target both 100:1
!
router ldp
!
interface lo
 ip address 127.0.0.1/8
 ip address 1.1.1.1/32 secondary
 ipv6 address ::1/128
!
interface eth0
 ip address dhcp
!
interface eth2
 ip vrf forwarding VRF1
 ip address 30.1.1.2/24
!
interface eth1
 ip address 40.1.1.1/24
 ipv6 address 3000::2/64
 label-switching
 enable-ldp ipv4
!
interface eth3
!
interface eth4
!
interface eth5
!
interface eth6
!
 exit
!
router ospf 100
 ospf router-id 1.1.1.1
 network 1.1.1.1/32 area 0.0.0.0
 network 40.1.1.0/24 area 0.0.0.0
!
router bgp 100
 bgp router-id 1.1.1.1
 neighbor 2.2.2.2 remote-as 100
 neighbor 2.2.2.2 update-source lo
 !
 address-family ipv4 unicast
 neighbor 2.2.2.2 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 2.2.2.2 activate
 exit-address-family
 !
 address-family ipv4 vrf VRF1
 aggregate-address 20.20.20.0/24 as-set summary-only
 neighbor 30.1.1.1 remote-as 200
 neighbor 30.1.1.1 activate
 neighbor 30.1.1.1 route-map map1 in
 exit-address-family
```

© 2026 IP Infusion Inc. Proprietary

```
 !
 exit
!
```

PE2

```
!
router-id 2.2.2.2
!
ip vrf management
!
ip vrf VRF1
 rd 200:1
 route-target both 100:1
!
router ldp
!
interface lo
 ip address 127.0.0.1/8
 ip address 2.2.2.2/32 secondary
 ipv6 address ::1/128
!
interface eth0
 ip address dhcp
!
interface eth1
 ip vrf forwarding VRF1
 ip address 50.1.1.2/24
!
interface eth2
 ip address 40.1.1.1/24
 ipv6 address 3000::1/64
 label-switching
 enable-ldp ipv4
!
interface eth3
!
interface eth4
!
interface eth5
!
interface eth6
!
 exit
!
router ospf 100
 ospf router-id 2.2.2.2
 network 2.2.2.2/32 area 0.0.0.0
 network 40.1.1.0/24 area 0.0.0.0
!
router bgp 100
 bgp router-id 2.2.2.2
 neighbor 1.1.1.1 remote-as 100
 neighbor 1.1.1.1 update-source lo
 !
 address-family ipv4 unicast
 neighbor 1.1.1.1 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 1.1.1.1 activate
 exit-address-family
 !
 address-family ipv4 vrf VRF1
 neighbor 50.1.1.1 remote-as 200
 neighbor 50.1.1.1 activate
```

© 2026 IP Infusion Inc. Proprietary

```
 exit-address-family
 !
 exit
!
```

CE1

```
interface eth1
 ip address 30.1.1.1/24
!
interface lo1
 ip address 20.20.20.1/32
!
interface lo2
 ip address 20.20.20.2/32
!
interface lo3
 ip address 20.20.20.3/32
!
nterface lo4
 ip address 20.20.20.4/32
!
 exit
!
router bgp 300
 neighbor 30.30.30.2 remote-as 100
 !
 address-family ipv4 unicast
 neighbor 30.30.30.2 activate
 redistribute connected
  exit-address-family
 !
 exit
!
```

Validation

```
PE1#sh ip bgp vpnv4 all summary
BGP router identifier 1.1.1.1, 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   Desc
2.2.2.2            4   100          8          7       1      0      0  00:02:10             0

Total number of neighbors 1

Total number of Established sessions 1

BGP VRF VRF1 Route Distinguisher: 100:1
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   Desc
30.1.1.1           4   200          6          7       1      0      0  00:02:10             0

Total number of neighbors 1

Total number of Established sessions 1


PE1#sh ip bgp vpnv4 all
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
Route Distinguisher: 100:1 (Default for VRF VRF1)
*> l 20.20.20.0/24    0.0.0.0               0        100        32768  ?        -
s>   20.20.20.2/32    30.1.1.1              0        100        0   3000 200 ?        -
s>   20.20.20.3/32    30.1.1.1              0        100        0   3000 200 ?        -
s>   20.20.20.4/32    30.1.1.1              0        100        0   3000 200 ?        -
s>   20.20.20.5/32    30.1.1.1              0        100        0   3000 200 ?        -
s>   20.20.20.6/32    30.1.1.1              0        100        0   3000 200 ?        -
s>   20.20.20.7/32    30.1.1.1              0        100        0   3000 200 ?        -
s>   20.20.20.8/32    30.1.1.1              0        100        0   3000 200 ?        -
s>   20.20.20.9/32    30.1.1.1              0        100        0   3000 200 ?        -
s>   20.20.20.20/32   30.1.1.1              0        100        0   3000 200 ?        -
 Announced routes count = 10
 Accepted routes count = 0
PE1#sh ip bgp neighbors 2.2.2.2 advertised-routes

For address family: VPNv4 Unicast
BGP table version is 3, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, a add-path, b back-up, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf   Weight Path   Ext-Color
*>i  20.20.20.0/24    1.1.1.1               0        100       32768  ?        -

Total number of prefixes 1
```

Configuration for Aggregate-Address in VRF Address-Family with IPv6

Topology

This sample topology provides basic connectivity for IPv6 Aggregated address.

Figure 98. BGP Aggregate-Address under Address-Family with IPv6

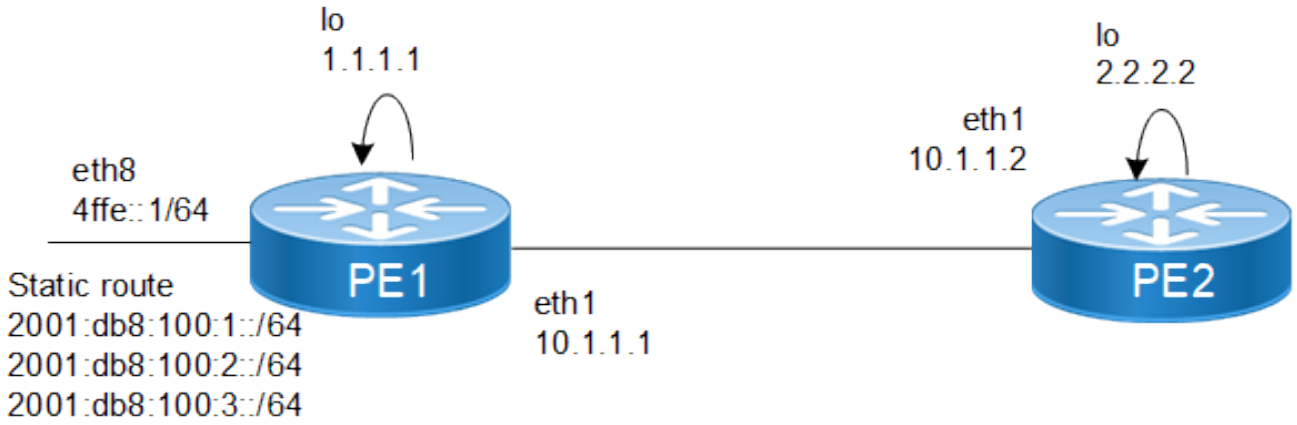

PE1

1. Configure Interfaces:

```
PE1(config)#interface eth1
PE1(config-if)#ip address 10.1.1.1/24
PE1(config-if)#ipv6 address 3000::1/64
PE1(config-if)#label-switching
PE1(config-if)#enable-ldp ipv4
PE1(config-if)#commit
PE1(config-if)#exit

PE1(config)#interface lo
PE1(config-if)#ip address 127.0.0.1/8
PE1(config-if)#ip address 1.1.1.1/32 secondary
PE1(config-if)#ipv6 address ::1/128
PE1(config-if)#commit
PE1(config-if)#exit
```

2. Configure OSPF (IPv4):

```
PE1(config)#router ospf 100
PE1(config-router)#network 1.1.1.1/32 area 0.0.0.0
PE1(config-router)#network 10.1.1.0/24 area 0.0.0.0
PE1(config-router)#commit
PE1(config-router)#exit
```

3. Configure BGP (IPv4 and IPv6):

```
PE1(config)#router bgp 100
PE1(config-router)#neighbor 2.2.2.2 remote-as 100
PE1(config-router)#neighbor 2.2.2.2 update-source lo
PE1(config-router)#address-family vpnv4 unicast
PE1(config-router-af)#neighbor 2.2.2.2 activate
PE1(config-router-af)#exit-address-family
PE1(config-router)#address-family vpnv6 unicast
PE1(config-router-af)#neighbor 2.2.2.2 activate
PE1(config-router-af)#exit-address-family
PE1(config-router)#address-family ipv6 vrf vrf1
PE1(config-router-af)#aggregate-address 2001:db8:100::/48 summary-only
PE1(config-router-af)#redistribute static
PE1(config-router-af)#exit-address-family
PE1(config-router)#commit
PE1(config-router)#exit
```

4. Configure Static IPv6 Routes in vrf1:

```
PE1(config)#ipv6 route vrf vrf1 2001:db8:100:1::/64 eth8
PE1(config)#ipv6 route vrf vrf1 2001:db8:100:2::/64 eth8
PE1(config)#ipv6 route vrf vrf1 2001:db8:100:3::/64 eth8
```

PE2

1. Configure Interfaces:

```
(config)#interface eth1
(config-if)#ip address 10.1.1.2/24
(config-if)#label-switching
(config-if)#enable-ldp ipv4
(config-if)#commit
(config-if)#exit

(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 2.2.2.2/32 secondary
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#exit
```

2. Configure OSPF (IPv4):

```
(config)#router ospf 100
```

```
(config-router)#network 2.2.2.2/32 area 0.0.0.0
(config-router)#network 10.1.1.0/24 area 0.0.0.0
```

3. Configure BGP (IPv4/IPv6/VPN):

```
(config)#router bgp 100
(config-router)#neighbor 1.1.1.1 remote-as 100
(config-router)#neighbor 1.1.1.1 update-source lo
(config-router-af)#address-family vpnv4 unicast
(config-router-af)#neighbor 1.1.1.1 activate
(config-router-af)#exit-address-family
(config-router)#address-family vpnv6 unicast
(config-router-af)#neighbor 1.1.1.1 activate
(config-router-af)#exit-address-family
(config-router)#address-family ipv6 vrf vrf1
(config-router-af)#exit-address-family
(config-router)#commit
(config-router)#exit
```

Running Configuration with IPv6

PE1

```
!
interface eth1
 ip address 10.1.1.1/24
 ipv6 address 3000::1/64
 label-switching
 enable-ldp ipv4
!
interface lo
 ip address 127.0.0.1/8
 ip address 1.1.1.1/32 secondary
 ipv6 address ::1/128
!
 exit
!
router ospf 100
 network 1.1.1.1/32 area 0.0.0.0
 network 10.1.1.0/24 area 0.0.0.0
!
router bgp 100
 neighbor 2.2.2.2 remote-as 100
 neighbor 2.2.2.2 update-source lo
 !
 address-family vpnv4 unicast
 neighbor 2.2.2.2 activate
 exit-address-family
 !
 address-family vpnv6 unicast
 neighbor 2.2.2.2 activate
 exit-address-family
 !
 address-family ipv6 vrf vrf1
 aggregate-address 2001:db8:100::/48 summary-only
 redistribute static
 exit-address-family
 !
 exit
!
ipv6 route vrf vrf1 2001:db8:100:1::/64 eth8
ipv6 route vrf vrf1 2001:db8:100:2::/64 eth8
ipv6 route vrf vrf1 2001:db8:100:3::/64 eth8
!
!
```

```
end
```

PE2

```
!
interface eth1
 ip address 10.1.1.2/24
 label-switching
 enable-ldp ipv4
!
interface lo
 ip address 127.0.0.1/8
 ip address 2.2.2.2/32 secondary
 ipv6 address ::1/128
!
 exit
!
router ospf 100
 network 2.2.2.2/32 area 0.0.0.0
 network 10.1.1.0/24 area 0.0.0.0
!
router bgp 100
 neighbor 1.1.1.1 remote-as 100
 neighbor 1.1.1.1 update-source lo
 !
 address-family vpnv4 unicast
 neighbor 1.1.1.1 activate
 exit-address-family
 !
 address-family vpnv6 unicast
 neighbor 1.1.1.1 activate
 exit-address-family
 !
 address-family ipv6 vrf vrf1
 exit-address-family
 !
 exit
!
!
end
```

Validation

PE1

```
PE1#show ip bgp vpnv6 all
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
Route Distinguisher: 100:1 (Default for VRF vrf1)
*> l 2001:db8:100::/48
                      ::                            0        100       32768  ?        -
s>   2001:db8:100:1::/64
                      ::                            0        100       32768  ?        -
s>   2001:db8:100:2::/64
                      ::                            0        100       32768  ?        -
s>   2001:db8:100:3::/64
                      ::                            0        100       32768  ?        -
 Announced routes count = 4
 Accepted routes count = 0
PE1#
```

PE2

```
PE2#show ip bgp vpnv6 all summary
BGP router identifier 10.16.42.104, local AS number 100
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   Desc
1.1.1.1            4   100         39         38       4      0      0  00:01:39             1

Total number of neighbors 1

Total number of Established sessions 1


OcNOS#sh ip bgp vpnv6 all
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
Route Distinguisher: 100:1
*>i  2001:db8:100:1::/64
                      ::ffff:1.1.1.1                0        100       0    ?        -
*>i  2001:db8:100:2::/64
                      ::ffff:1.1.1.1                0        100       0    ?        -
*>i  2001:db8:100:3::/64
                      ::ffff:1.1.1.1                0        100       0    ?        -
 Announced routes count = 0
 Accepted routes count = 3


After aggregation
==========================
PE2#sh ip bgp vpnv6 all
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
Route Distinguisher: 100:1
*>i  2001:db8:100::/48
                      ::ffff:1.1.1.1                0        100       0    ?        -
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 200:1 (Default for VRF vrf1)
*>i  2001:db8:100::/48
                      ::ffff:1.1.1.1                0        100       0    ?        -
 Announced routes count = 0
 Accepted routes count = 1
PE2#
```

BGP RPKI-Based Route Validation

Overview

Resource Public Key Infrastructure (RPKI) is a security framework designed to mitigate the risk of BGP prefix hijacking by cryptographically verifying that an Autonomous System (AS) is authorized to announce a given IP prefix.

In OcNOS, RPKI-based BGP Origin Validation allows the router to download Route Origin Authorizations (ROAs) from an RPKI server via the RTR protocol. The downloaded ROAs are then used to validate incoming BGP routes, ensuring that only legitimate prefixes are considered during best path selection.

This feature improves routing security by reducing the acceptance and propagation of invalid routes.

Feature Characteristics

- **ROA Retrieval:** Supports downloading ROAs from multiple (up to 10) RPKI servers over TCP or SSH transport.
- **Per-AF and Per-VRF Support:** Validation can be enabled on a per-address-family (IPv4/IPv6 unicast) and per-VRF basis.
- **Validation States:** Each route is tagged with one of the three validation state:
  - **Valid (V):** Prefix-AS match found in ROA.
  - **Invalid (I):** Prefix-AS mismatch or not authorized.
  - **Not-Found (N):** No corresponding ROA.
- **Flexible Policy Control:** Route-map support for matching on RPKI state (valid, invalid, not-found) to set attributes such as local preference.
- **Best Path Selection Control:**
  - Option to consider only valid/not-found routes for path selection.
  - Configurable to allow invalid routes in best path preference.
- **Dynamic Updates:** ROA updates are applied in real time from RPKI servers.

Benefits

- **Enhanced Security:** Prevents acceptance of hijacked or misconfigured routes.
- **Operational Flexibility:** Operators can tune route selection with route-maps or allow invalid routes for troubleshooting.
- **Standards Compliance:** Implements BGP Origin Validation as per RPKI-based validation standards.
- **Granular Control:** Policies can be applied per Address Family (AF)or Virtual Routing and Forwarding (VRF), giving operators flexibility in deploying validation gradually.
- **Improved Resilience:** Reduces propagation of invalid prefixes across the Internet routing system.

© 2026 IP Infusion Inc. Proprietary

Configuration

This section describes the configuration procedure for enabling RPKI-based BGP origin validation.

Prerequisites

- BGP is already configured (router bgp (ASN)) on the device.
- You have reachable RPKI RTR server IP(s) and credentials (if using SSH).
- Make sure the required transport ports (TCP or SSH) are reachable between the OcNOS and the RPKI server(s).
- Decide per address-family (IPv4/IPv6) and per VRF where you want validation turned on.

Topology

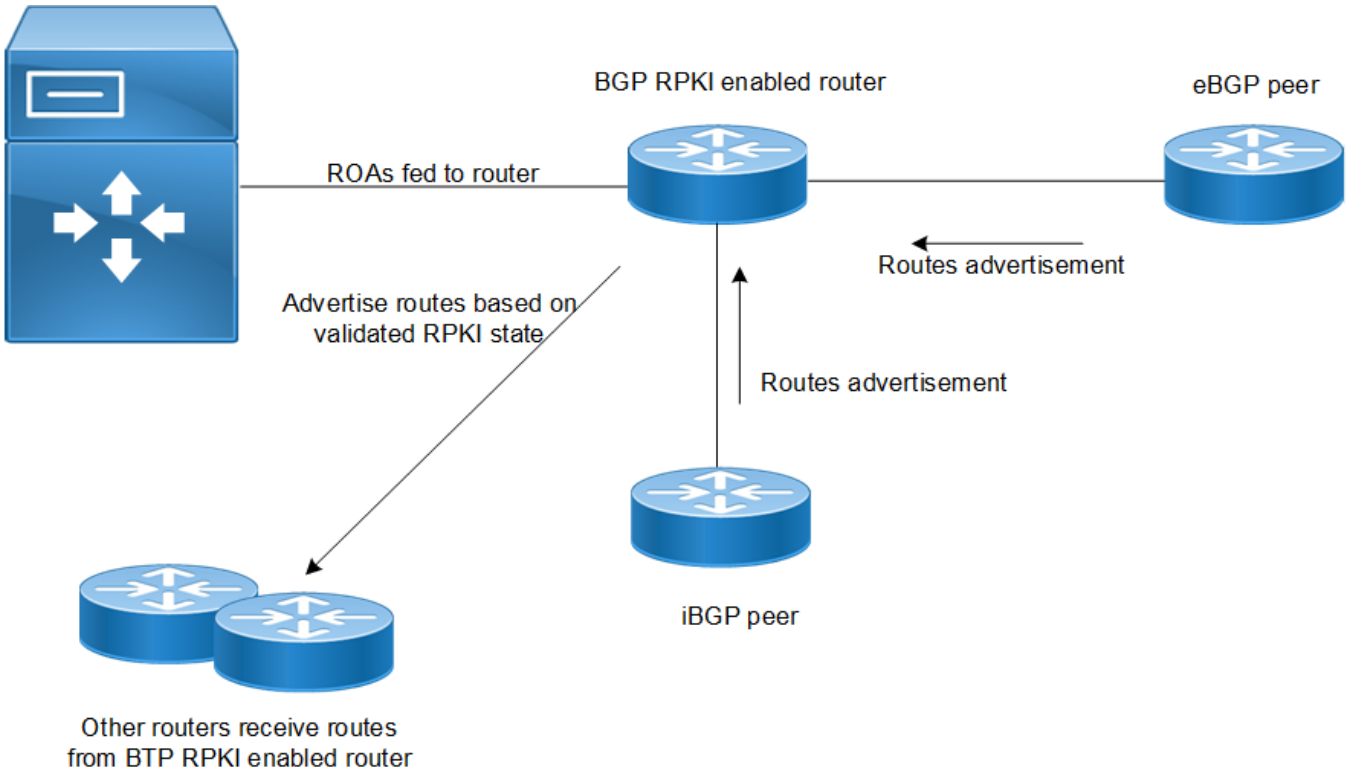

In this topology, the RPKI Validator stores and distributes Route Origin Authorizations (ROAs), which contain information about the prefixes and its authorized originating AS numbers. The validator communicates these validated ROAs to the BGP RPKI-enabled router.

The BGP RPKI-enabled router establishes a session with the RPKI Validator using either TCP or SSH, depending on the configuration. Upon receiving ROAs, the router validates the BGP route advertisements against the authorized prefix–origin pairs.

Based on the validation results, the router marks each route as Valid, Invalid, or NotFound, and applies routing policies accordingly.

Only routes that pass the validation check (Valid) are used for forwarding or are advertised to other peers.

- ROAs fed to router: The RPKI Validator sends validated prefix-origin data to the BGP RPKI-enabled router.

© 2026 IP Infusion Inc. Proprietary

- Advertise routes based on validated RPKI state: The router advertises only validated routes to its BGP peers.
- iBGP and eBGP peers: Both internal (iBGP) and external (eBGP) peers receive route updates from the RPKI-enabled router.
- Other routers: Routers within the same AS or network domain receive routes that have already been validated, ensuring route authenticity and preventing prefix hijacking.

1. Configure one or more RPKI servers to establish RTR sessions and download Route Origin Authorizations (ROAs).

```
ocnos(config)# router bgp 100
ocnos(config-router)# bgp rpki server 10.30.0.85 ssh user test encrypt 0 password test refresh 1
retry 1 expire 600
ocnos(config-router)# bgp rpki server 192.168.1.233 tcp refresh 1 retry 1 expire 600
ocnos(config-router)# commit
```

2. Enable origin validation on the required address family (AF) or VRF. This allows BGP routes to be tagged with an RPKI validation state:
    - V: Valid
    - I: Invalid
    - N: Not-found

```
ocnos(config)# router bgp 100
ocnos(config-router)# address-family ipv4 unicast
ocnos(config-router-af)# bgp origin-as validation-enable
ocnos(config-router-af)# commit
```

3. Configure BGP to consider RPKI validation state in the best-path selection process. Invalid routes are excluded, and preference is given in the following order: valid > not-found (unless modified by policy).

```
ocnos(config-router-af)# bgp origin-as bestpath use-validity
ocnos(config-router-af)# commit
```

4. Permit invalid routes to participate in best-path selection but assign them the lowest preference.

```
ocnos(config-router-af)# bgp origin-as bestpath allow-invalid
ocnos(config-router-af)# commit
```

5. Use route-maps to define policy actions, such as setting local preference, based on the RPKI validation state of a route.

    Example of the route map:

```
route-map RPKI-1 permit 3
 match rpki valid
 set local-preference 100

route-map RPKI-1 permit 5
 match rpki not-found
 set local-preference 200

route-map RPKI-1 permit 10
 match rpki invalid
 set local-preference 300
```

```
ocnos(config-router)# address-family ipv4 unicast
ocnos(config-router-af)# neighbor 100.1.1.2 activate
ocnos(config-router-af)# neighbor 100.1.1.2 route-map RPKI-1 in
ocnos(config-router-af)# commit
```

6. Remove or rollback RPKI configuration:
    - Disable Validation in an AF/VRF:

© 2026 IP Infusion Inc. Proprietary

```
ocnos(config-router-af)# no bgp origin-as validation-enable
ocnos(config-router-af)# commit
```

- Remove an RPKI Server:

```
ocnos(config)# router bgp 100
ocnos(config-router)# no bgp rpki server 10.30.0.85
ocnos(config-router)# commit
```

Validation

Verify server session for the following:

State: Established and Synced: TRUE in show bgp rpki server detail.

```
OCNOS#show bgp rpki server detail
BGP RPKI Server Information
 Server Address: 155.155.1.1:
    Transport: TCP:3323
    RTR Version: 1
    State: established
    Synced: TRUE
    Uptime: 00:00:39
    ROAs (IPv4/IPv6): 9/3
    Configured Refresh-Interval: 15 seconds
    Configured Retry-Interval: 10 seconds
    Configured Expire-Interval: 600 seconds
    Actual Refresh-Interval: 5 seconds
    Actual Retry-Interval: 5 seconds
    Actual Expire-Interval: 600 seconds
    Rest of time to Refresh-Interval expiration: 2 seconds
    Rest of time to Expire-Interval expiration: 597 seconds
    ToBeDeleted: FALSE
```

Verify validation state on routes

```
ocnos# show bgp origin-as validity ipv4
BGP table version is 28, local router ID is 1.1.1.1

Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,

              l - labeled, S Stale, x-EVPN

Origin codes: i - IGP, e - EGP, ? - incomplete

Description : Ext-Color - Extended community color



Origin-AS validation codes: V valid, I invalid, N not-found, D disabled



    Network          Next Hop            Metric    LocPrf   Weight Path   Ext-Color

N*>   1.2.0.0/16       0.0.0.0              0        100       32768  ?        -

I*>   1.2.11.0/24      0.0.0.0              0        100       32768  i        -

I*                     100.1.1.7            0        100       32768  ?        -

N*    1.2.21.30/32     0.0.0.0              0        100       32768  i        -

N*>   1.6.0.0/16       0.0.0.0              0        100       32768 {300,9583} ?        -
```

```
I*>i  1.6.14.0/24      100.1.1.7           0       100       0    ?        -
I*                     100.1.1.7           20      100       32768  ?        -
V*>i  1.6.136.0/24     100.4.1.5           0       100       0   9583 ?        -
V*                     100.2.1.3           0       100       0   300 9583 ?        -
N*    2.0.0.1/32       0.0.0.0             0       100       32768  i        -
N*>i  2.2.2.2/32       100.1.1.2           0       100       0    ?        -
I*    2.3.4.5/32       0.0.0.0             0       100       32768  i        -
I*>   3.3.3.3/32       100.2.1.3           0       100       0   300 ?        -
I*    3.4.5.6/32       0.0.0.0             0       100       32768  i        -
N*>i  4.4.4.4/32       100.4.1.4           0       100       0   400 ?        -
N*                     100.2.1.3           0       100       0   300 400 ?        -
I*>i  5.5.5.5/32       100.4.1.5           0       100       0   9583 ?        -
I*                     100.2.1.3           0       100       0   300 9583 ?        -
N*>i  7.7.7.7/32       100.1.1.7           0       100       0    ?        -
N*                     100.1.1.7           20      100       32768  ?        -
I*>   8.8.8.8/32       100.1.1.8           0       100       0   400 1300 ?        -
N*    33.44.55.66/32   0.0.0.0             0       100       32768  i        -
N*>   100.1.1.0/24     100.1.1.8           0       100       0   400 1300 ?        -
N* i                   100.1.1.2           0       100       0    ?        -
N* i                   100.1.1.7           0       100       0    ?        -
N*                     0.0.0.0             1       100       32768  ?        -
I*>   100.2.1.0/24     100.2.1.3           0       100       0   300 ?        -
I*                     100.1.1.8           0       100       0   400 1300 ?        -
I* i                   100.1.1.7           0       100       0    ?        -
N*>   100.3.1.0/24     100.2.1.3           0       100       0   300 ?        -
I*>   100.4.1.0/24     100.2.1.3           0       100       0   300 400 ?        -
I* i                   100.1.1.2           0       100       0    ?        -
N*>   172.16.181.0/24  100.2.1.3           0       100       0   300 ?        -
N*                     100.1.1.8           0       100       0   400 1300 ?        -
N* i                   100.1.1.2           0       100       0    ?        -
N* i                   100.1.1.7           0       100       0    ?        -


Total number of prefixes 21
```

Implementation Examples

Incomplete Global RPKI Adoption

Not all Internet Service Providers (ISPs) or regional networks fully participate in RPKI validation.

Example Scenario: Peers from regions with limited RPKI coverage advertise valid prefixes that appear as invalid due to missing ROAs.

Action: It may permit invalid routes from trusted peers or specific regions to maintain global connectivity.

Route Leak or Failover Scenarios

During failover or traffic engineering events, alternate paths may temporarily appear as invalid.

Example Scenario: A backup eBGP link originates a prefix from a different AS path than specified in the ROA.

Action:The route may be allowed conditionally (for example, through route maps) to ensure reachability during the transition period.

Commands

The BGP RPKI as origin validation feature introduces the following commands:

bgp rpki server

Use this command to configure an RPKI cache server using either TCP or SSH transport, and to set the port and timer parameters (refresh, retry, expire), along with authentication details when SSH is used.

Use the **no** parameter of this command to remove an existing RPKI server configuration from the BGP instance.

Command Syntax

```
bgp rpki server (A.B.C.D or X:X::X:X) (tcp|) (port (port number)) (refresh (1 - 86400 )) (retry (1 -
7200 )) (expire (600 - 17200))

bgp rpki server (A.B.C.D or X:X::X:X) ssh user (user name) encrypt (0|1) password (PASSWORD) (port
(PORT NUMBER)) (refresh (1 - 86400 )) (retry (1 - 7200 )) (expire (600 - 17200))
```

Parameters

A.B.C.D
Specifies the IPv4 address of the RPKI cache server.

X:X::X:X
Specifies the IPv6 address of the RPKI cache server.

tcp
Specifies TCP as the communication protocol between the router and the RPKI server.

ssh
Specifies SSH as the communication protocol between the router and the RPKI server for secure communication.

user
Defines the username used to authenticate with the RPKI server.

encrypt
Specifies encrypted or not encrypted for a password

0
Defines unencrypted password (key)

1
Defines encrypted password (key)

password
Defines the BGP encrypted password (key) up to maximum 218 characters for an ssh connection.

port
Specifies the TCP or SSH port number to connect to the RPKI server.

refresh
Specifies the time interval, in seconds, to refresh the cache from the RPKI server.

retry
Specifies the time interval, in seconds, to retry the connection to the RPKI server if the previous attempt fails.

expire
Specifies the cache expiration interval, in seconds, after which cached RPKI data is discarded if not refreshed.

Default

None

© 2026 IP Infusion Inc. Proprietary

Command Mode

Router mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example illustrates how to specify rpki server:

```
OcNOS(config-router)# bgp rpki server 10.30.0.85 tcp port 3323 refresh 600 retry 120 expire 7200

OcNOS(config-router)# no bgp rpki server 10.30.0.85

OcNOS(config-router)#bgp rpki server 1.1.1.1 ssh user test encrypt 0 password 123
OcNOS(config-router)#commit
```

© 2026 IP Infusion Inc. Proprietary

bgp origin-as validation-enable

Use this command to enable BGP Origin-AS (AS Origin) Validation using RPKI. When enabled, the router validates the origin AS of received BGP prefixes against the ROA information downloaded from RPKI servers..

Use the **no** parameter of this command to disable BGP Origin-AS validation.

Command Syntax

```
bgp origin-as validation-enable
no bgp origin-as validation-enable
```

Parameters

None

Default

None

Command Mode

Address Family Configuration Mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example enables Origin-AS validation for BGP bestpath selection in the current address family configuration:

```
OcNOS(config)#router bgp 100
OcNOS(config-router)#address-family ipv4 unicast
OcNOS(config-router-af)# bgp origin-as validation-enable
```

The following example disables Origin-AS validation for BGP bestpath selection in the current address family configuration:

```
OcNOS(config-router-af)# no bgp origin-as validation-enable
```

© 2026 IP Infusion Inc. Proprietary

bgp origin-as bestpath

Use this command to control how BGP selects the best path when RPKI Origin-AS validation is enabled.

This command determines whether BGP considers RPKI validation results during the bestpath selection process or allows paths with an invalid RPKI state to be selected.

Use the **no** form of this command to restore the default bestpath behavior.

Command Syntax

```
bgp origin-as bestpath (allow-invalid | use validity)
no bgp origin-as bestpath (allow-invalid | use validity)
```

Parameters

allow-invalid
Enables to handle a route with invalid RPKI state for the best path selection

Use-validity
Enables to use origin-as validation for the bestpath selection

Default

None

Command Mode

Address Family Configuration Mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example enables Origin-AS validation for BGP bestpath selection in the current address family configuration:

```
OcNOS(config)#router bgp 100
OcNOS(config-router)#address-family ipv4 unicast
OcNOS(config-router-af)# bgp origin-as bestpath allow-invalid
```

The following example disables Origin-AS validation for BGP bestpath selection in the current address family configuration:

```
OcNOS(config-router-af)# no bgp origin-as bestpath allow-invalid
```

© 2026 IP Infusion Inc. Proprietary

match rpki

Use this command to match BGP routes based on their Resource Public Key Infrastructure (RPKI) validation status in a route-map.

Use the **no** parameter of this command to remove an existing RPKI match configuration from the route-map.

Command Syntax

```
match rpki {valid | invalid | not-found}
no match rpki {valid | invalid | not-found}
```

Parameters

valid
Matches routes that have a valid RPKI validation status.

invalid
Matches routes that have an invalid RPKI validation status.

not-found
Matches routes whose RPKI validation status is unknown (not found).

Default

None

Command Mode

Route map mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example illustrates to match only invalid RPKI routes in a route-map and to remove the invalid RPKI routes.

```
ocnos#config  terminal
ocnos(config)#route-map 1
ocnos(config-route-map)# match rpki invalid

ocnos(config-route-map)# no match rpki valid
```

© 2026 IP Infusion Inc. Proprietary

show bgp rpki table ipv4

Use this command to display the IPv4 RPKI ROA table.

The command shows all validated ROA entries downloaded from configured RPKI servers.

Command Syntax

```
show bgp rpki table ipv4 (A.B.C.D/M ((covered|matched) (as-no <1-4294967295>)
```

Parameters

A.B.C.D/M
Displays ROA entries that include or relate to the specified IPv4 prefix.

covered
Displays ROA entries where the specified prefix is covered by a larger ROA prefix.

matched
Filters the ROA table to display entries associated with a specific authorized origin AS number.

as-no <1–4294967295>
Filters the ROA table to display entries associated with a specific authorized origin AS number.

Default

None

Command Mode

Execution mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example displays IPv4 RPKI ROA table:

```
ocnos#show bgp rpki table ipv4
 BGP RPKI, ROA list
 1.0.0.0/24                                Maxlen:24 AS:13335 Server:10.30.0.85
                                           Maxlen:24 AS:13335 Server:192.168.1.233
 1.0.64.0/18                               Maxlen:18 AS:18144 Server:10.30.0.85
                                           Maxlen:18 AS:18144 Server:192.168.1.233
 1.1.1.0/24                                Maxlen:24 AS:13335 Server:10.30.0.85
                                           Maxlen:24 AS:13335 Server:192.168.1.233
 1.1.4.0/22                                Maxlen:22 AS:4134 Server:10.30.0.85
                                           Maxlen:22 AS:4134 Server:192.168.1.233
 1.1.16.0/20                               Maxlen:20 AS:4134 Server:10.30.0.85
                                           Maxlen:20 AS:4134 Server:192.168.1.233
 1.2.9.0/24                                Maxlen:24 AS:4134 Server:10.30.0.85
                                           Maxlen:24 AS:4134 Server:192.168.1.233
 1.2.10.0/24                               Maxlen:24 AS:4134 Server:10.30.0.85
                                           Maxlen:24 AS:4134 Server:192.168.1.233
```

Explanation of output fields:

© 2026 IP Infusion Inc. Proprietary

Field	Description
Prefix	Shows the IPv4 prefix (ROA prefix) published by the RPKI trust anchor.
Maxlen	Indicates the maximum prefix length allowed by the ROA.
AS	Displays the authorized Origin AS number for the prefix, as specified in the ROA.
Server	Indicates the RPKI server (RPKI RTR server address) from which the ROA entry was received.

© 2026 IP Infusion Inc. Proprietary

show bgp rpki table ipv6

Use this command to display the IPv6 RPKI ROA table.

The command shows all validated ROA entries downloaded from configured RPKI servers.

Command Syntax

```
show bgp rpki table ipv6 (X:X::X:X/M ((covered|matched) (as-no <1-4294967295>|)|)|)
```

Parameters

X:X::X:X/M
Displays ROA entries that include or relate to the specified IPv6 prefix.

covered
Displays ROA entries where the specified prefix is covered by a larger ROA prefix.

matched
Filters the ROA table to display entries associated with a specific authorized origin AS number.

as-no <1–4294967295>
Filters the ROA table to display entries associated with a specific authorized origin AS number.

Default

None

Command Mode

Execution mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example displays IPv6 RPKI ROA table:

```
ocnos#show bgp rpki table ipv6
 BGP RPKI, ROA list
 2001:200::/32                          Maxlen:32 AS:2500 Server:192.168.1.233
 2001:200:136::/48                      Maxlen:48 AS:9367 Server:192.168.1.233
 2001:200:1ba::/48                      Maxlen:48 AS:24047 Server:192.168.1.233
 2001:200:900::/40                      Maxlen:40 AS:7660 Server:192.168.1.233
 2001:200:e00::/40                      Maxlen:40 AS:4690 Server:192.168.1.233
 2001:200:8000::/35                     Maxlen:35 AS:4690 Server:192.168.1.233
 2001:200:c000::/35                     Maxlen:35 AS:23634 Server:192.168.1.233
 2001:200:e000::/35                     Maxlen:35 AS:7660 Server:192.168.1.233
 2001:218::/32                          Maxlen:32 AS:2914 Server:192.168.1.233
 2001:218:2000:2::/64                   Maxlen:64 AS:4058 Server:192.168.1.233
 2001:218:2000:11::/64                  Maxlen:64 AS:55569 Server:192.168.1.233
 2001:218:2000:21::/64                  Maxlen:64 AS:55569 Server:192.168.1.233
 2001:218:2002::/48                     Maxlen:48 AS:2914 Server:192.168.1.233
```

Explanation of output fields:

Field	Description
Prefix	Shows the IPv6 prefix (ROA prefix) published by the RPKI trust anchor.
Maxlen	Indicates the maximum prefix length allowed by the ROA.
AS	Displays the authorized Origin AS number for the prefix, as specified in the ROA.
Server	Indicates the RPKI server (RPKI RTR server address) from which the ROA entry was received.

© 2026 IP Infusion Inc. Proprietary

show bgp origin-as validity ipv4

Use this command to display the RPKI Origin-AS validation state of IPv4 BGP routes.The command shows whether each route is classified as valid, not-found, or invalid based on the ROA information received from RPKI servers.

Command Syntax

```
show bgp origin-as validity ipv4 (valid|not-found|invalid|) (vrf WORD|)
```

Parameters

valid
Displays only the IPv4 prefixes that passed RPKI Origin-AS validation.

not-found
Displays routes for which no corresponding ROA entry exists.

invalid
Displays routes that fail RPKI Origin-AS validation.

vrf WORD
Displays RPKI validation results for the specified VRF instance.

Default

None

Command Mode

Execution mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example displays bgp origin-as validity for IPv4 address:

```
ocnos#show bgp origin-as validity ipv4
BGP table version is 9, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
              l - labeled, S Stale, x-EVPN MPLS
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

Origin-AS validation codes: V valid, I invalid, N not-found, D disabled

    Network          Next Hop            Metric    LocPrf   Weight Path   Ext-Color
V*>   1.6.136.0/24     100.2.1.3              0        100        0   300 9583 ?        -
N*>i  4.4.4.4/32       100.4.1.4              0        200        0   400 ?        -
N*                     100.2.1.3              0        200        0   300 400 ?        -
I*>   5.5.5.5/32       100.2.1.3              0        300        0   300 9583 ?        -
N*>   100.3.1.0/24     100.2.1.3              0        200        0   300 9583 ?        -
I*>   100.4.1.0/24     100.2.1.3              0        300        0   300 9583 ?        -
N*>   172.16.181.0/24  100.2.1.3              0        200        0   300 9583 ?        -

Total number of prefixes 6
```

show bgp origin-as validity ipv6

Use this command to display the RPKI Origin-AS validation state of IPv6 BGP routes.The command shows whether each route is classified as valid, not-found, or invalid based on the ROA information received from RPKI servers.

Command Syntax

```
show bgp origin-as validity ipv6 (valid|not-found|invalid|) (vrf WORD|)
```

Parameters

valid
Displays only the IPv6 prefixes that passed RPKI Origin-AS validation.

not-found
Displays routes for which no corresponding ROA entry exists.

invalid
Displays routes that fail RPKI Origin-AS validation.

vrf WORD
Displays RPKI validation results for the specified VRF instance.

Default

None

Command Mode

Execution mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example displays origin-as validity for IPv6 address:

```
ocnos#show bgp origin-as validity ipv6
BGP table version is 270, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
              l - labeled, S Stale, x-EVPN MPLS
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

Origin-AS validation codes: V valid, I invalid, N not-found, D disabled

    Network          Next Hop                    Metric    LocPrf     Weight Path
N*>i  2001::100:1:1:0/120
                      2001::100:1:1:2(fe80::20c:29ff:feea:6a30)
                                                  0        100       0    ?        -
N*>   2001::100:2:1:0/120
                      2001::100:2:1:3(fe80::20c:29ff:fedf:6e11)
                                                  0        100       0   300 ?        -
N*>   2001::100:3:1:0/120
                      2001::100:2:1:3(fe80::20c:29ff:fedf:6e11)
                                                  0        100       0   300 ?        -
N* i                   2001::100:4:1:4              0        100       0   400 ?        -
N*>i  2001::100:4:1:0/120
```

```
                     2001::100:1:1:2(fe80::20c:29ff:feea:6a30)
                                                0        100       0    ?        -
N*                    2001::100:2:1:3(fe80::20c:29ff:fedf:6e11)
                                                0        100       0   300 400 ?        -
N*>i  2601:647:6300:8dc0::/64
                     2001::100:1:1:2(fe80::20c:29ff:feea:6a30)
                                                0        100       0    ?        -

Total number of prefixes 5
```

show bgp rpki server

Use this command to display the configured RPKI servers, their transport protocol, operational state, uptime, and the number of received Route Origin Authorizations (ROAs) for both IPv4 and IPv6.

Command Syntax

```
show bgp rpki server  (detail|summary)
```

Parameters

None

Default

None

Command Mode

Privileged execution mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example displays rpki server show command:

```
ocnos#show bgp rpki server
 Server Address                          Transport   State         Uptime      ROAs (IPv4/IPv6)
 10.16.99.108                            TCP:3323    established   00:02:27    9/3
 155.155.1.1                             SSH:2222    established   00:02:26    9/3

ocnos#show bgp rpki server detail
BGP RPKI Server Information
 Server Address: 155.155.1.1:
    Transport: TCP:3323
    RTR Version: 1
    State: established
    Synced: TRUE
    Uptime: 00:00:39
    ROAs (IPv4/IPv6): 9/3
    Configured Refresh-Interval: 15 seconds
    Configured Retry-Interval: 10 seconds
    Configured Expire-Interval: 600 seconds
    Actual Refresh-Interval: 5 seconds
    Actual Retry-Interval: 5 seconds
    Actual Expire-Interval: 600 seconds
    Rest of time to Refresh-Interval expiration: 2 seconds
    Rest of time to Expire-Interval expiration: 597 seconds
    ToBeDeleted: FALSE
```

Explanation of output fields:

Field	Description
Server Address	IP address of the configured RPKI server.
Transport	Transport protocol and port used to connect to the RPKI

Field	Description
	server (for example: TCP:or SSH:)
RTR Version	Version of the RPKI-to-Router (RTR) protocol negotiated with the server.
State	Current operational state of the RPKI server
Synced	Indicates whether the router has successfully synchronized ROA data with the server.
Uptime	Duration for which the server connection has been active or the current status since the last state change.
ROAs (IPv4/IPv6)	Number of valid Route Origin Authorizations received from the server, separated by IPv4 and IPv6 counts.
Configured Refresh-Interval	User-configured refresh timer for requesting updated validation data.
Configured Retry-Interval	User-configured retry interval for reconnecting after a session failure.
Configured Expire-Interval	User-configured maximum validity period for cached ROA data.
Actual Refresh-Interval	Refresh interval currently in use, as negotiated with the server.
Actual Retry-Interval	Retry interval currently in use, based on server negotiation.
Actual Expire-Interval	Expire interval currently in use, based on protocol negotiation.
Rest of time to Refresh-Interval expiration	Remaining time until the next refresh request is triggered.
Rest of time to Expire-Interval expiration	Remaining time before cached validation data becomes invalid.
ToBeDeleted	Indicates whether the server entry is marked for deletion.

Glossary

Key Terms/Acronym	Description
Border Gateway Protocol (BGP)	The standardized exterior gateway protocol used to exchange routing information between autonomous systems (ASes) on the Internet.
Resource Public Key Infrastructure (RPKI)	A framework designed to secure the Internet's routing infrastructure by cryptographically verifying that an AS is authorized to originate a specific IP prefix.
Origin Validation	A process where a BGP router validates the origin AS of a received route against RPKI data to determine if the route is legitimate.

© 2026 IP Infusion Inc. Proprietary

Key Terms/Acronym	Description
Route Origin Authorization (ROA)	digitally signed object that specifies which AS is authorized to announce a particular IP prefix.

Recursive Next-hop Resolution with MPLS Next-hop

Overview

Recursive Next-hop Resolution with MPLS Next-hop enables OcNOS to resolve recursive BGP and Static routes using MPLS transport paths instead of IP next-hops. This feature allows seamless packet forwarding through an MPLS core where BGP is not active, supporting scalable BGP-Free Core deployments. The system recognizes available MPLS Label Switched Paths (LSPs) as valid forwarding options, ensuring reachability across label-switched domains without requiring additional routing protocols in the backbone.

Feature Characteristics

- Supports route resolution over LDP, RSVP-TE, and Segment Routing (SR-MPLS) tunnels.
- Automatically updates route resolution when the underlying MPLS transport changes.
- Integrates with MPLS ECMP, FRR, and TI-LFA for resiliency and load balancing.
- Displays MPLS-resolved recursive routes in standard routing table views.
- Maintains consistent behavior for both BGP and Static recursive routes.
- Operates transparently without additional configuration changes to the routing process.

Benefits

Enables BGP-Free Core Operation: Backbone routers forward traffic using MPLS without running BGP.

Simplified Architecture: Reduces control-plane load and configuration complexity.

Improved Resiliency: Leverages MPLS protection mechanisms for faster failover.

Optimized Traffic Flow: Supports equal-cost load sharing over multiple MPLS paths.

Operational Visibility: Provides clear differentiation of MPLS-resolved routes for easier troubleshooting.

Limitation

- BGP multipath is not supported for BGP-Free Core deployments.
- If multiple BGP paths exist for the same prefix, the RIB selects only one path for MPLS next-hop resolution.
- Equal-Cost Multipath (ECMP) and Fast Reroute (FRR) are provided only by the underlying MPLS transport for that single resolved path.

Configuration

This section describes the Recursive Next-hop Resolution with MPLS Next-hop feature, illustrating its operation with a sample topology, configuration examples, and validation in an MPLS LDP and SR-MPLS deployments.

Configuration with MPLS LDP

Topology (MPLS LDP)

In this topology:

- PE1 or PE2 are Provider Edge routers connected to customer devices (CE1 or CE2).

© 2026 IP Infusion Inc. Proprietary

- P1 or P2 are core Provider routers.
- LDP is enabled on all PE and P routers to exchange MPLS labels and establish LSPs.
- OSPF and IS-IS are used as the IGP to advertise reachability between LDP peers.
- BGP is used between PE routers for VPN route exchange (VRFs).

Figure 99. Recursive Next-hop Resolution with MPLS Next-hop in MPLS LDP Deployment

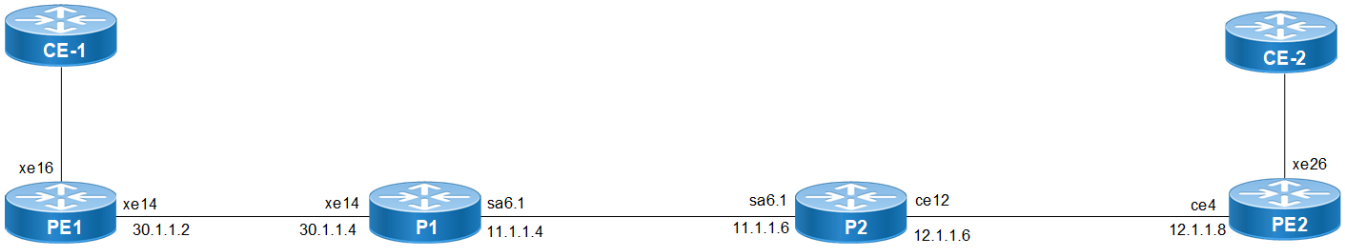

PE1: 2.2.2.2 (loopback), xe14 to P1, xe16.80 to CE1.
P1: 4.4.4.4 (loopback), sa6.1 to P2, xe14 to PE1.
P2: 6.6.6.6 (loopback), sa6.1 to P1, ce12 to PE2.
PE2: 8.8.8.8 (loopback), ce4 to P2, xe26.80 to CE2.

Enable MPLS LDP on Core and PE Routers

- LDP requires a router ID (typically loopback).
- LDP is enabled per interface connecting MPLS-capable peers.
- OSPF or IS-IS should be running to provide IGP reachability for LDP session establishment.

```
router ldp
 router-id 2.2.2.2
!
!
interface lo
 ip address 2.2.2.2/32 secondary
!
interface xe14
 load-interval 30
 ip address 30.1.1.2/24
 label-switching
 ip ospf network point-to-point
 enable-ldp ipv4
!
interface xe16.80
 encapsulation dot1q 80
 ip address 80.1.1.1/24
 ip ospf network point-to-point
!
```

Configure IGP (OSPF or IS-IS)

- IGP ensures all LDP peers can discover each other.
- BFD is optional for fast detection of failures.

```
!
router ospf 1
 ospf router-id 2.2.2.2
 bfd all-interfaces
 network 2.2.2.2/32 area 0.0.0.0
 network 30.1.1.0/24 area 0.0.0.0
!
```

© 2026 IP Infusion Inc. Proprietary

Configure BGP (for VPNs or PE-to-PE routes)

- BGP exchanges prefixes between PE routers.
- MPLS labels are used to forward traffic between PE routers over the LDP-established LSPs.

```
!
router bgp 64000
 bgp router-id 2.2.2.2
  neighbor 8.8.8.8 remote-as 64000
 neighbor 8.8.8.8 update-source lo
 !
 address-family ipv4 unicast
 redistribute connected
  neighbor 8.8.8.8 activate
 exit-address-family
 !
!
```

Validation (MPLS LDP)

Verify LDP Sessions

- **Expected output:** Session state = OPERATIONAL
- Confirms label exchange between LDP peers (e.g., PE1 ↔ P1).

```
PE1#show mpls ldp session
! [execution timestamp : 2025 Oct 09 16:25:19]
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime  AutoCreate
       4.4.4.4            xe14       Passive   OPERATIONAL   30    00:40:43      No
```

Verify IP Routes

- Check that all networks (PE, P, CE) are reachable.
- Routes learned via OSPF or ISI-S or BGP should point to next-hop over MPLS.

```
PE1#show ip route
! [execution timestamp : 2025 Oct 09 16:25:08]
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, installed 1d14h08m, last update 1d14h08m ago
O            4.4.4.4/32 [110/2] via 30.1.1.4, xe14, installed 00:33:59, last update 00:33:59 ago
O            6.6.6.6/32 [110/3] via 30.1.1.4, xe14, installed 00:33:59, last update 00:33:59 ago
O            8.8.8.8/32 [110/4] via 30.1.1.4, xe14, installed 00:33:50, last update 00:33:50 ago
O            11.1.1.0/24 [110/2] via 30.1.1.4, xe14, installed 00:33:59, last update 00:33:59 ago
O            12.1.1.0/24 [110/3] via 30.1.1.4, xe14, installed 00:33:59, last update 00:33:59 ago
O            26.10.1.0/24 [110/4] via 30.1.1.4, xe14, installed 00:33:50, last update 00:33:50 ago
C            30.1.1.0/24 is directly connected, xe14, installed 00:40:53, last update 00:40:53 ago
C            80.1.1.0/24 is directly connected, xe16.80, installed 1d14h08m, last update 1d14h08m ago
B            81.1.1.0/24 [200/0] via 8.8.8.8 (recursive via 30.1.1.4), installed 00:33:41, last
update 00:33:41 ago
C            127.0.0.0/8 is directly connected, lo, installed 1d14h09m, last update 1d14h09m ago
```

© 2026 IP Infusion Inc. Proprietary

```
O           222.1.1.0/24 [110/4] via 30.1.1.4, xe14, installed 00:33:50, last update 00:33:50 ago

Gateway of last resort is not set
```

Verify MPLS Labels (ILM or FTN)

- ILM table shows labels received from peers.
- FTN table shows outgoing labels and interfaces for forwarding.

```
PE1#show mpls ilm-table
! [execution timestamp : 2025 Oct 09 16:25:11]
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP , K - CLI , V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy,       U - unknown, UPStr - upstream

ILM-ECMP LDP: Enabled, SR: Enabled
Code    FEC/VRF/L2CKT    ILM-ID      In-Label    Out-Label   In-Intf    Out-
Intf/VRF       Nexthop                pri  Algo-Num  UpTime    UPStr peers
   B>   vrf100             2           26305       Nolabel     N/A        N/A              N/A
             Yes  N/A       00:40:49
   B>   vrf200             1           26304       Nolabel     N/A        N/A              N/A
             Yes  N/A       00:40:49


PE1#show mpls forwarding-table
! [execution timestamp : 2025 Oct 09 16:25:13]
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Enabled, SR: Enabled
Code    FEC                FTN-ID    Nhlfe-ID  Tunnel-ID   Pri   Out-Label    Out-
Intf    ELC       Nexthop         Algo-Num   UpTime
   L>   4.4.4.4/32          1         7         -           -     -            -          -
         -               N/A        00:02:04
                                      4         -
           Yes   0            xe14        Yes       30.1.1.4        -          -
   L>   6.6.6.6/32          2         10        -           -     -            -          -
         -               N/A        00:02:04
                                      9         -
           Yes   24960        xe14        No        30.1.1.4        -          -
   L>   8.8.8.8/32          3         12        -           -     -            -          -
         -               N/A        00:02:04
                                      11        -
           Yes   24961        xe14        No        30.1.1.4        -          -
   L>   11.1.1.0/24         4         7         -           -     -            -          -
         -               N/A        00:40:37
                                      4         -
           Yes   0            xe14        Yes       30.1.1.4        -          -
   L>   12.1.1.0/24         5         15        -           -     -            -          -
         -               N/A        00:40:37
                                      14        -
           Yes   24962        xe14        No        30.1.1.4        -          -
   L>   26.10.1.0/24        6         17        -           -     -            -          -
         -               N/A        00:40:37
                                      16        -
           Yes   24963        xe14        No        30.1.1.4        -          -
   L>   222.1.1.0/24        10        36        -           -     -            -          -
         -               N/A        00:33:54
                                      29        -
           Yes   24965        xe14        No        30.1.1.4        -          -
```

Example from PE1:

```
FEC          Out-Label   Out-Intf   Nexthop
6.6.6.6/32   24960       xe14       30.1.1.4
8.8.8.8/32   24961       xe14       30.1.1.4
```

Verify MPLS Recursive Route Resolution on PE1

The `show ip route database` and `show ip route mpls` outputs confirm that PE1 resolves BGP routes via MPLS Label Switched Paths (LSPs) established through LDP.

- The BGP route `81.1.1.0/24` (customer network behind PE2) is learned from the BGP peer `8.8.8.8` (PE2).
- The next hop `8.8.8.8` is not directly connected; it is recursively resolved via the IGP route to `30.1.1.4` (toward P1).
- The additional line (`recursive via MPLS FTN-ID 3, nhlfe-ix 12`) indicates that PE1 uses an MPLS forwarding entry (FTN-ID 3) for this BGP route, proving that traffic destined for the remote CE network (`81.1.1.0/24`) is label-switched through the MPLS core rather than IP-routed.

```
PE1#show ip route database
! [execution timestamp : 2025 Oct 09 16:25:26]
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    *>  2.2.2.2/32 is directly connected, lo, installed 1d14h08m, last update 1d14h08m ago
O        2.2.2.2/32 [110/1] is directly connected, lo, installed 00:34:27, last update 00:34:27 ago
O    *>  4.4.4.4/32 [110/2] via 30.1.1.4, xe14, installed 00:34:17, last update 00:34:17 ago
O    *>  6.6.6.6/32 [110/3] via 30.1.1.4, xe14, installed 00:34:17, last update 00:34:17 ago
O    *>  8.8.8.8/32 [110/4] via 30.1.1.4, xe14, installed 00:34:08, last update 00:34:08 ago
B        8.8.8.8/32 [200/0] via 8.8.8.8 inactive, installed 00:33:59, last update 00:33:59 ago
O    *>  11.1.1.0/24 [110/2] via 30.1.1.4, xe14, installed 00:34:17, last update 00:34:17 ago
O    *>  12.1.1.0/24 [110/3] via 30.1.1.4, xe14, installed 00:34:17, last update 00:34:17 ago
B        12.1.1.0/24 [200/0] via 8.8.8.8 (recursive via 30.1.1.4), installed 00:33:59, last update
00:33:59 ago
O    *>  26.10.1.0/24 [110/4] via 30.1.1.4, xe14, installed 00:34:08, last update 00:34:08 ago
B        26.10.1.0/24 [200/0] via 8.8.8.8 (recursive via 30.1.1.4), installed 00:33:59, last update
00:33:59 ago
C    *>  30.1.1.0/24 is directly connected, xe14, installed 00:41:11, last update 00:41:11 ago
O        30.1.1.0/24 [110/1] is directly connected, xe14, installed 00:34:27, last update 00:34:27
ago
C    *>  80.1.1.0/24 is directly connected, xe16.80, installed 1d14h08m, last update 1d14h08m ago
O        80.1.1.0/24 [110/1] is directly connected, xe16.80, installed 1d14h08m, last update 1d14h08m
ago
B    *>  81.1.1.0/24 [200/0] via 8.8.8.8 (recursive via 30.1.1.4), installed 00:33:59, last update
00:33:59 ago
     *>              [200/0] via 8.8.8.8 (recursive via MPLS FTN-ID 3, nhlfe-ix 12)
C    *>  127.0.0.0/8 is directly connected, lo, installed 1d14h09m, last update 1d14h09m ago
O    *>  222.1.1.0/24 [110/4] via 30.1.1.4, xe14, installed 00:34:08, last update 00:34:08 ago
B        222.1.1.0/24 [200/0] via 8.8.8.8 (recursive via 30.1.1.4), installed 00:33:59, last update
00:33:59 ago

Total number of IPv4 routes 20

Gateway of last resort is not set
```

```
PE1#show ip route mpls
```

```
! [execution timestamp : 2025 Oct 09 16:25:31]
IP Route Table for VRF "default"
B            81.1.1.0/24 [200/0] via 8.8.8.8 (recursive via MPLS FTN-ID 3, nhlfe-ix 12), installed
00:34:04, last update 00:34:04 ago

Gateway of last resort is not set
```

Test End-to-End Reachability (MPLS LDP)

- Ping or traceroute should reach CE2 via MPLS LSPs (PE1 → P1 → P2 → PE2).
- MPLS labels are swapped at core nodes (P1, P2) automatically.

```
PE1#ping 81.1.1.1
Press CTRL+C to exit
PING 81.1.1.1 (81.1.1.1) 100(128) bytes of data.
108 bytes from 81.1.1.1: icmp_seq=1 ttl=64 time=0.709 ms
108 bytes from 81.1.1.1: icmp_seq=2 ttl=64 time=0.583 ms
108 bytes from 81.1.1.1: icmp_seq=3 ttl=64 time=0.564 ms

--- 81.1.1.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2059ms
rtt min/avg/max/mdev = 0.564/0.618/0.709/0.064 ms
PE1#
PE1#
PE1#traceroute 81.1.1.1
traceroute to 81.1.1.1 (81.1.1.1), 30 hops max, 60 byte packets
 1  81.1.1.1 (81.1.1.1)  0.901 ms  0.852 ms  0.981 ms
```

Observed Behavior (MPLS LDP)

- **LDP Sessions:** PE1 establishes LDP with P1, P1 with P2, and P2 with PE2.
- **MPLS Forwarding:** Labels assigned to each destination allow PE1 to reach PE2's CE network (`81.1.1.0/24`) via core LSP.
- **BGP and MPLS Interaction:** BGP routes for customer networks are advertised and attached to labels. PE1 forwards traffic using MPLS labels without IP routing lookup at core routers.
- **Validation confirms:**
  - IP routes are correct (`show ip route`).
  - MPLS label tables are populated (`show mpls ilm-table and show mpls forwarding-table`).
  - End-to-end connectivity works (`ping and traceroute`).

Configuration with Segment Routing MPLS (SR-MPLS)

Topology (SR-MPLS)

In this topology:

- PE1 or PE2 are Provider Edge routers connected to customer devices (CE1 or CE2).
- P1 or P2 are core Provider routers.
- SR-MPLS is enabled on PE and P routers to provide MPLS transport using IGP-advertised Prefix-SIDs (no LDP sessions are required).
- OSPF is used as the IGP to advertise reachability and distribute Prefix-SID mappings across the network.
- BGP is enabled between PE routers to exchange customer prefixes; recursive next-hop resolution uses the SR-MPLS transport path for forwarding.

© 2026 IP Infusion Inc. Proprietary

Figure 100. Recursive Next-hop Resolution with MPLS Next-hop in SR-MPLS Deployment

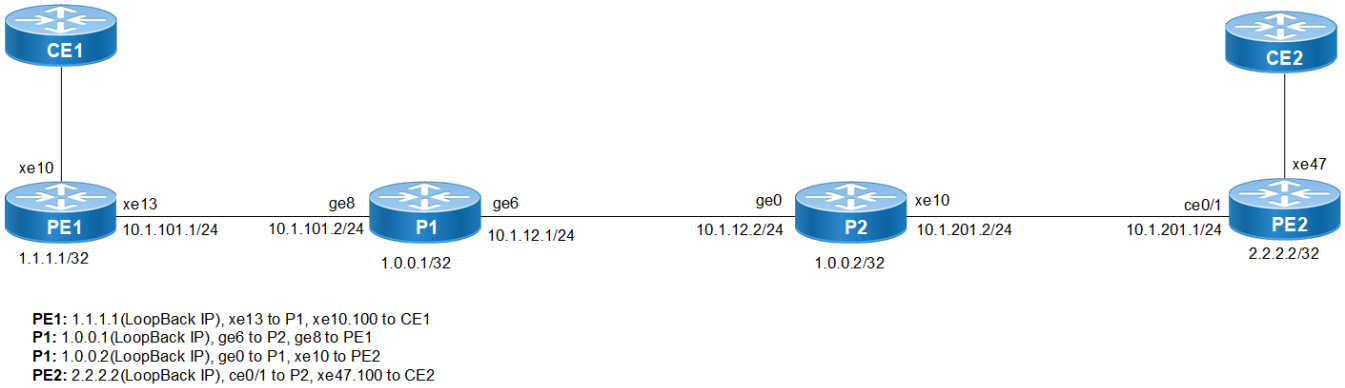

Configure Loopback and Prefix-SID

Each router must have a loopback router-id and a Prefix-SID used by SR-MPLS to build labeled transport paths.

- `prefix-sid index` assigns the SID index for the loopback prefix.
- `no-php` disables penultimate hop popping for the Prefix-SID label.

```
!
interface lo
 ip address 127.0.0.1/8
 ip address 1.0.0.1/32 secondary
 ipv6 address ::1/128
 prefix-sid index 101 no-php
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
```

Enable Label Switching on Core-facing Interfaces

Enable MPLS label switching on interfaces that participate in the SR-MPLS transport. This ensures SR labels are pushed or swapped or popped through the MPLS transport path.

P1

```
!
interface ge6
 load-interval 30
 ip address 10.1.12.1/24
 mtu 9216
 label-switching
 ip router isis 100
!
interface ge7
!
interface ge8
 load-interval 30
 ip address 10.1.101.2/24
 mtu 9216
 label-switching
 ip router isis 100
!
```

Configure IGP with SR enabled (OSPF)

OSPF provides reachability and distributes Prefix-SID mappings required for SR-MPLS transport.

© 2026 IP Infusion Inc. Proprietary

- `segment-routing mpls` command enables SR-MPLS in the OSPF process.
- (Optional) `fast-reroute per-prefix ti-lfa` command enables fast reroute using TI-LFA.

```
!
router ospf 100
 ospf router-id 1.0.0.1
 network 1.0.0.1/32 area 0.0.0.0
 network 10.1.12.0/24 area 0.0.0.0
 network 10.1.101.0/24 area 0.0.0.0
 segment-routing mpls
 fast-reroute per-prefix ti-lfa area 0.0.0.0
!
```

Configure BGP (PE-to-PE routes)

BGP exchanges customer prefixes between PE routers. With this feature, recursive BGP next-hops can be resolved via SR-MPLS transport.

```
!
router bgp 100
 bgp router-id 1.1.1.1
 neighbor 2.2.2.2 remote-as 100
 neighbor 2.2.2.2 update-source lo
 neighbor 2.2.2.2 advertisement-interval 0
 !
 address-family ipv4 unicast
 redistribute connected
 neighbor 2.2.2.2 activate
 exit-address-family
 !
 exit
!
```

Validation (SR-MPLS)

Verify Segment Routing State

- **Expected output:** `SR State: SR_ENABLED` and `Operational state: enabled`
- Confirms Segment Routing is enabled for the OSPF process and SR-MPLS transport is operational.

```
PE1#show ip ospf segment-routing state

OSPF process 100 Segment-Routing:
SR State: SR_ENABLED
SRGB Start: 16000,  SRGB Range: 8000
SRLB Start: 14080,  SRLB Range: 1920
Operational state: enabled

PE2#show ip ospf segment-routing state

OSPF process 100 Segment-Routing:
SR State: SR_ENABLED
SRGB Start: 16000,  SRGB Range: 8000
SRLB Start: 14080,  SRLB Range: 1920
Operational state: enabled
```

Verify Prefix-SID Mapping

- **Expected output:** Mapping table lists loopback prefixes with SID indexes for all nodes.
- Confirms Prefix-SIDs are advertised and learned across the SR domain.

```
PE1#show ip ospf segment-routing mapping-table active
OSPF process ID: 100
```

© 2026 IP Infusion Inc. Proprietary

```
Conflict Resolution Policy: Quarantine

Prefix              SID Index      Range       Flags
1.0.0.1/32          101            1
1.0.0.2/32          102            1
1.1.1.1/32          111            1
2.2.2.2/32          222            1

Number of mapping entries in Active Table: 4


PE2#show ip ospf segment-routing mapping-table active
OSPF process ID: 100
Conflict Resolution Policy: Quarantine

Prefix              SID Index      Range       Flags
1.0.0.1/32          101            1
1.0.0.2/32          102            1
1.1.1.1/32          111            1
2.2.2.2/32          222            1

Number of mapping entries in Active Table: 4
```

Verify IP Routes

Check that all networks (PE, P, CE) are reachable. Routes learned via OSPF or BGP must provide reachability to the remote PE loopback and remote customer subnet.

```
PE1#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            1.0.0.1/32 [110/2] via 10.1.101.2, xe13, installed 00:30:00, last update 00:30:00 ago
O            1.0.0.2/32 [110/3] via 10.1.101.2, xe13, installed 00:25:48, last update 00:25:48 ago
C            1.1.1.1/32 is directly connected, lo, installed 00:37:24, last update 00:37:24 ago
O            2.2.2.2/32 [110/4] via 10.1.101.2, xe13, installed 00:21:35, last update 00:21:35 ago
O            10.1.12.0/24 [110/2] via 10.1.101.2, xe13, installed 00:30:00, last update 00:30:00 ago
C            10.1.101.0/24 is directly connected, xe13, installed 00:33:18, last update 00:33:18 ago
O            10.1.201.0/24 [110/3] via 10.1.101.2, xe13, installed 00:22:22, last update 00:22:22 ago
C            100.100.1.0/24 is directly connected, xe10.100, installed 00:33:18, last update 00:33:18
ago
C            127.0.0.0/8 is directly connected, lo, installed 00:55:01, last update 00:55:01 ago
B            200.200.1.0/24 [200/0] via 2.2.2.2 (recursive via 10.1.101.2), installed 00:21:29, last
update 00:21:29 ago

Gateway of last resort is not set

PE2#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            1.0.0.1/32 [110/3] via 10.1.201.2, ce0/1, installed 00:30:44, last update 00:30:44 ago
O            1.0.0.2/32 [110/2] via 10.1.201.2, ce0/1, installed 00:30:44, last update 00:30:44 ago
O            1.1.1.1/32 [110/4] via 10.1.201.2, ce0/1, installed 00:30:44, last update 00:30:44 ago
```

```
C            2.2.2.2/32 is directly connected, lo, installed 00:31:30, last update 00:31:30 ago
O            10.1.12.0/24 [110/2] via 10.1.201.2, ce0/1, installed 00:30:44, last update 00:30:44 ago
O            10.1.101.0/24 [110/3] via 10.1.201.2, ce0/1, installed 00:30:44, last update 00:30:44
ago
C            10.1.201.0/24 is directly connected, ce0/1, installed 00:31:30, last update 00:31:30 ago
B            100.100.1.0/24 [200/0] via 1.1.1.1 (recursive via 10.1.201.2), installed 00:30:39, last
update 00:30:39 ago
C            127.0.0.0/8 is directly connected, lo, installed 00:59:47, last update 00:59:47 ago
C            200.200.1.0/24 is directly connected, xe47.100, installed 00:31:30, last update 00:31:30
ago

Gateway of last resort is not set
```

Verify MPLS Labels (ILM or FTN) for SR-MPLS

- ILM table shows incoming label entries learned for SR-OSPF.
- FTN table shows outgoing labels and interfaces used to forward labeled traffic.

```
PE1#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP , K - CLI , V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy,       U - unknown, UPStr - upstream

ILM-ECMP LDP: Disabled, SR: Enabled
Code    FEC/VRF/L2CKT    ILM-ID      In-Label    Out-Label   In-Intf    Out-
Intf/VRF       Nexthop                 pri  Algo-Num  UpTime    UPStr peers
   O>   2.2.2.2/32         5           16222       16222       N/A        xe13             10.1.101.2
             Yes  0         00:21:50
   O>   1.0.0.2/32         4           16102       16102       N/A        xe13             10.1.101.2
             Yes  0         00:26:03
   O>   1.0.0.1/32         3           16101       16101       N/A        xe13             10.1.101.2
             Yes  0         00:30:15
   O>   1.1.1.1/32         1           16111       Nolabel     N/A        N/A              127.0.0.1
             Yes  0         00:06:29
   O>   10.1.101.2/32      2           24960       3           N/A        xe13             10.1.101.2
             Yes  0         00:30:17

PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Enabled
Code    FEC                FTN-ID    Nhlfe-ID  Tunnel-ID   Pri   Out-Label    Out-
Intf    ELC       Nexthop         Algo-Num   UpTime
   O>   1.0.0.1/32          1         6         -           -     -            -           -
         -               0          00:30:19
                                      5         0           Yes   16101        xe13        No
10.1.101.2      -          -
   O>   1.0.0.2/32          2         9         -           -     -            -           -
         -               0          00:26:07
                                      8         0           Yes   16102        xe13        No
10.1.101.2      -          -
   O>   2.2.2.2/32          3         12        -           -     -            -           -
         -               0          00:21:54
                                      11        0           Yes   16222        xe13        No
10.1.101.2      -          -

PE2#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP , K - CLI , V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
```

```
       P - SR Policy,       U - unknown, UPStr - upstream

ILM-ECMP LDP: Disabled, SR: Enabled
Code    FEC/VRF/L2CKT    ILM-ID      In-Label    Out-Label   In-Intf    Out-
Intf/VRF       Nexthop                pri  Algo-Num  UpTime    UPStr peers
   O>   2.2.2.2/32         1           16222       Nolabel     N/A        N/A              127.0.0.1
             Yes  0         00:31:53
   O>   1.0.0.2/32         4           16102       16102       N/A        ce0/1            10.1.201.2
             Yes  0         00:31:08
   O>   1.0.0.1/32         3           16101       16101       N/A        ce0/1            10.1.201.2
             Yes  0         00:31:08
   O>   1.1.1.1/32         5           16111       16111       N/A        ce0/1            10.1.201.2
             Yes  0         00:15:45
   O>   10.1.201.2/32      2           24960       3           N/A        ce0/1            10.1.201.2
             Yes  0         00:31:46

PE2#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Enabled
Code    FEC                 FTN-ID    Nhlfe-ID  Tunnel-ID   Pri   Out-Label    Out-
Intf    ELC       Nexthop         Algo-Num   UpTime
   O>   1.0.0.1/32          1         4         -           -     -            -           -
         -               0          00:31:11
                                      3         0           Yes   16101        ce0/1       No
10.1.201.2      -          -
   O>   1.0.0.2/32           2         7         -           -     -            -           -
         -               0          00:31:11
                                      6         0           Yes   16102        ce0/1       No
10.1.201.2      -          -
   O>   1.1.1.1/32           3         16        -           -     -            -           -
         -               0          00:15:48
                                      15        0           Yes   16111        ce0/1       No
10.1.201.2      -          -
```

Verify MPLS Recursive Route Resolution on PE Routers

The `show ip route database` and `show ip route mpls` outputs confirm that PE routers resolve BGP routes via SR-MPLS transport.

On PE1

- The BGP route `200.200.1.0/24` (remote network behind PE2) is learned from the BGP peer `2.2.2.2` (PE2 loopback).
- The BGP next-hop `2.2.2.2` is recursively reachable through OSPF via `10.1.101.2` (toward P1).
- The additional line (`recursive via MPLS FTN-ID 3, nhlfe-ix 12`) confirms that the route is installed using an MPLS forwarding entry, proving SR-MPLS label forwarding through the core.

```
PE1#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"
O    *>   1.0.0.1/32 [110/2] via 10.1.101.2, xe13, installed 00:30:04, last update 00:30:04 ago
```

```
O    *>   1.0.0.2/32 [110/3] via 10.1.101.2, xe13, installed 00:25:52, last update 00:25:52 ago
C    *>   1.1.1.1/32 is directly connected, lo, installed 00:37:28, last update 00:37:28 ago
O         1.1.1.1/32 [110/1] is directly connected, lo, installed 00:32:55, last update 00:32:55 ago
O    *>   2.2.2.2/32 [110/4] via 10.1.101.2, xe13, installed 00:21:39, last update 00:21:39 ago
B         2.2.2.2/32 [200/0] via 2.2.2.2 inactive, installed 00:21:33, last update 00:21:33 ago
O    *>   10.1.12.0/24 [110/2] via 10.1.101.2, xe13, installed 00:30:04, last update 00:30:04 ago
C    *>   10.1.101.0/24 is directly connected, xe13, installed 00:33:22, last update 00:33:22 ago
O         10.1.101.0/24 [110/1] is directly connected, xe13, installed 00:32:55, last update 00:32:55
ago
O    *>   10.1.201.0/24 [110/3] via 10.1.101.2, xe13, installed 00:22:26, last update 00:22:26 ago
B         10.1.201.0/24 [200/0] via 2.2.2.2 (recursive via 10.1.101.2), installed 00:21:33, last
update 00:21:33 ago
C    *>   100.100.1.0/24 is directly connected, xe10.100, installed 00:33:22, last update 00:33:22
ago
C    *>   127.0.0.0/8 is directly connected, lo, installed 00:55:05, last update 00:55:05 ago
B    *>   200.200.1.0/24 [200/0] via 2.2.2.2 (recursive via 10.1.101.2), installed 00:21:33, last
update 00:21:33 ago
     *>                  [200/0] via 2.2.2.2 (recursive via MPLS FTN-ID 3, nhlfe-ix 12)

Total number of IPv4 routes 15

Gateway of last resort is not set

PE1#show ip route mpls
IP Route Table for VRF "default"
B            200.200.1.0/24 [200/0] via 2.2.2.2 (recursive via MPLS FTN-ID 3, nhlfe-ix 12), installed
00:21:38, last update 00:21:38 ago

Gateway of last resort is not set
```

On PE2

- The BGP route `100.100.1.0/24` (remote network behind PE1) is learned from the BGP peer `1.1.1.1` (PE1 loopback).
- The BGP next-hop `1.1.1.1` is recursively reachable through OSPF via `10.1.201.2` (toward P2).
- The additional line (`recursive via MPLS FTN-ID 3, nhlfe-ix 16`) confirms MPLS forwarding entry selection for SR-MPLS transport.

```
PE2#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"
O    *>   1.0.0.1/32 [110/3] via 10.1.201.2, ce0/1, installed 00:30:56, last update 00:30:56 ago
O    *>   1.0.0.2/32 [110/2] via 10.1.201.2, ce0/1, installed 00:30:56, last update 00:30:56 ago
O    *>   1.1.1.1/32 [110/4] via 10.1.201.2, ce0/1, installed 00:30:56, last update 00:30:56 ago
B         1.1.1.1/32 [200/0] via 1.1.1.1 inactive, installed 00:30:51, last update 00:30:51 ago
C    *>   2.2.2.2/32 is directly connected, lo, installed 00:31:42, last update 00:31:42 ago
O         2.2.2.2/32 [110/1] is directly connected, lo, installed 00:31:42, last update 00:31:42 ago
O    *>   10.1.12.0/24 [110/2] via 10.1.201.2, ce0/1, installed 00:30:56, last update 00:30:56 ago
O    *>   10.1.101.0/24 [110/3] via 10.1.201.2, ce0/1, installed 00:30:56, last update 00:30:56 ago
B         10.1.101.0/24 [200/0] via 1.1.1.1 (recursive via 10.1.201.2), installed 00:30:51, last
update 00:30:51 ago
C    *>   10.1.201.0/24 is directly connected, ce0/1, installed 00:31:42, last update 00:31:42 ago
O         10.1.201.0/24 [110/1] is directly connected, ce0/1, installed 00:31:42, last update
00:31:42 ago
B    *>   100.100.1.0/24 [200/0] via 1.1.1.1 (recursive via 10.1.201.2), installed 00:30:51, last
update 00:30:51 ago
     *>                  [200/0] via 1.1.1.1 (recursive via MPLS FTN-ID 3, nhlfe-ix 16)
```

```
C    *>   127.0.0.0/8 is directly connected, lo, installed 00:59:59, last update 00:59:59 ago
C    *>   200.200.1.0/24 is directly connected, xe47.100, installed 00:31:42, last update 00:31:42
ago

Total number of IPv4 routes 15

Gateway of last resort is not set

PE2#show ip route  mpls
IP Route Table for VRF "default"
B            100.100.1.0/24 [200/0] via 1.1.1.1 (recursive via MPLS FTN-ID 3, nhlfe-ix 16), installed
00:30:58, last update 00:30:58 ago

Gateway of last resort is not set
```

Test End-to-End Reachability (SR-MPLS)

Ping or traceroute must reach the remote customer subnet through SR-MPLS transport (labels are imposed at ingress and forwarded across the core).

From PE1 to CE2

```
PE1#ping 200.200.1.1
Press CTRL+C to exit
PING 200.200.1.1 (200.200.1.1) 100(128) bytes of data.
108 bytes from 200.200.1.1: icmp_seq=1 ttl=64 time=0.715 ms
108 bytes from 200.200.1.1: icmp_seq=2 ttl=64 time=0.665 ms
108 bytes from 200.200.1.1: icmp_seq=3 ttl=64 time=0.602 ms

--- 200.200.1.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2073ms
rtt min/avg/max/mdev = 0.602/0.660/0.715/0.046 ms

PE1#traceroute 200.200.1.1
traceroute to 200.200.1.1 (200.200.1.1), 30 hops max, 60 byte packets
 1  200.200.1.1 (200.200.1.1)  0.898 ms  0.874 ms  0.901 ms
```

From PE2 to CE1

```
PE2#ping 100.100.1.1
Press CTRL+C to exit
PING 100.100.1.1 (100.100.1.1) 100(128) bytes of data.
108 bytes from 100.100.1.1: icmp_seq=1 ttl=64 time=0.677 ms
108 bytes from 100.100.1.1: icmp_seq=2 ttl=64 time=0.642 ms
108 bytes from 100.100.1.1: icmp_seq=3 ttl=64 time=0.595 ms
108 bytes from 100.100.1.1: icmp_seq=4 ttl=64 time=0.617 ms

--- 100.100.1.1 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3104ms
rtt min/avg/max/mdev = 0.595/0.632/0.677/0.030 ms

PE2#traceroute 100.100.1.1
traceroute to 100.100.1.1 (100.100.1.1), 30 hops max, 60 byte packets
 1  100.100.1.1 (100.100.1.1)  0.782 ms  0.881 ms  0.915 ms
```

Observed Behavior (SR-MPLS)

- **Segment Routing State: PE1** and **PE2** report `SR_ENABLED` under OSPF, confirming SR-MPLS is active and operational in the transport domain.
- **Prefix-SID Distribution: PE1 (1.1.1.1/32, SID 111)**, **P1 (1.0.0.1/32, SID 101)**, **P2 (1.0.0.2/32, SID 102)**, and **PE2 (2.2.2.2/32, SID 222)** are present in the SR mapping table, confirming end-to-end Prefix-SID learning

across the core.

- **MPLS Forwarding (SR Labels):** SR-OSPF label entries are installed on **PE1** and **PE2**, and labels are used for transport forwarding across **P1** and **P2** (no LDP sessions are required).
- **BGP and MPLS Interaction: PE1** learns the remote customer prefix **200.200.1.0/24** from **PE2 (2.2.2.2)** and resolves it through SR-MPLS transport (`recursive via MPLS FTN-ID 3, nhlfe-ix 12`). Similarly, **PE2** learns **100.100.1.0/24** from **PE1 (1.1.1.1)** and resolves it through SR-MPLS (`recursive via MPLS FTN-ID 3, nhlfe-ix 16`).
- **Validation confirms:**
  - IP routes are correct (**PE1/PE2** `show ip route`).
  - SR state and Prefix-SIDs are present (**PE1/PE2** `show ip ospf segment-routing state` and `mapping-table active`).
  - MPLS SR label tables are populated (**PE1/PE2** `show mpls ilm-table` and `show mpls forwarding-table`).
  - End-to-end connectivity works (`ping` and `traceroute`).

Implementation Example

Use Case: Simplified MPLS Core for BGP-Free Operation

In a large network backbone, the service provider's core is designed for label-switched forwarding using MPLS. Dynamic routing, such as BGP, is enabled only at the network edges where service routes are originated or terminated. The core routers focus solely on MPLS transport, without maintaining full BGP tables.

With Recursive Next-Hop Resolution, OcNOS edge or aggregation routers can automatically resolve BGP or static routes through available MPLS paths, even when the next-hop is not directly reachable via IP. This allows traffic to traverse the MPLS core seamlessly, enabling a scalable and operationally efficient BGP-free core design. The feature also leverages MPLS protection mechanisms for rapid failover and improved resiliency within the core.

Revised Commands

show ip route mpls

Added the `mpls` and `mpls summary` parameters to the `show ip route` command to display routing table entries resolved through MPLS next-hops (NHLFE). For more details, refer to the `show ip route` command section in the *OcNOS System Management Guide*.

show ip route database

Added the `MPLS FTN-ID <ID>, nhlfe-ix <INDEX>` field to the `show ip route database` command output. This field specifies the MPLS next-hop used for recursive resolution. For more details, refer to the `show ip route database` command section in the *OcNOS System Management Guide*.

show rib txlist

The show rib txlist (page 1234) command includes the fields (`req_mpls, new_req_mpls, fib_notify_mpls, seq_mpls`) to display MPLS-specific route installation and synchronization details between RIB and FIB.

© 2026 IP Infusion Inc. Proprietary

Troubleshooting

The following steps can be used to validate recursive next-hop resolution through MPLS paths:

1. Verify that RIB has MPLS transport information: Use the diagnostic command `show rib mpls ipv4` (hidden debug command) to confirm that MPLS transport details are received and available for resolution.
2. Verify that recursive routes are installed: Check the routing table entries resolved via MPLS next-hops using the show commands `show ip route (|summary | database)` and `show ip route mpls (|summary)`.
3. Verify recursive route installation in hardware: Confirm that MPLS-resolved prefixes are programmed in the forwarding tables using the following show commands:
    - `show hsl prefix-table ipv4 fib-id 0`
    - For MPLS transport validation:
        - LDP or SR-MPLS tunnels: `show hsl mpls ftn nhlfe-ix <index>`
        - RSVP-TE tunnels: `show hsl mpls tunnel tunnel-id <id>`

© 2026 IP Infusion Inc. Proprietary

show rib txlist

Use this command to display internal Routing Information Base (RIB) transaction lists that are pending or currently being transmitted to the Forwarding Information Base (FIB). This command helps monitor routes awaiting installation or synchronization between the routing and forwarding planes. It is mainly used for debugging and operational verification of route processing states.

Command Syntax

```
show rib (txlist | pending-txlist) (ipv4 | ipv6) (vrf WORD|)
```

Parameters

txlist
Displays the list of routes that are currently being transmitted from RIB to FIB.

pending-txlist
Displays the list of routes queued for transmission to the FIB.

ipv4
Displays information for IPv4 routes.

ipv6
Displays information for IPv6 routes.

vrf WORD
Specifies the VRF name for which the RIB transaction information is displayed.

Default

None

Command Mode

Execution mode and Privileged execution mode

Applicability

Introduced before OcNOS version 1.3.

Added the fields (`req_mpls, new_req_mpls, fib_notify_mpls, seq_mpls`) in OcNOS version 7.0.0.

Example

The fields (`req_mpls, new_req_mpls, fib_notify_mpls, seq_mpls`) show MPLS-specific route installation and synchronization details between RIB and FIB.

```
#show rib txlist ipv4
    IPv4:
    0x02739a6c: 127.0.0.0/8        req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
    0x02810cdc: 1.2.0.0/24         req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(5)
    0x0281071c: 1.1.1.1/32         req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(5)
    0x028d9c5c: 2.2.2.2/32         req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(4)
    0x028da42c: 11.11.11.11/32     req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
    0x028dab2c: 12.12.12.12/32     req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
```

© 2026 IP Infusion Inc. Proprietary

```
    0x028daf7c: 13.13.13.13/32     req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
    0x028db0cc: 9.9.9.9/32         req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
    0x028db32c: 10.10.10.10/32     req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
    MARKER MPLS*
    MARKER*

#show rib pending-txlist ipv4 vrf VRF3
IPv4 pending:
0x02634680: 192.0.5.0/24       req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x02634830: 192.0.6.0/24       req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x026349e0: 192.0.7.0/24       req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x02634bb0: 192.0.8.0/24       req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x02652490: 192.0.9.0/24       req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x02652660: 192.0.10.0/24      req(1) new_req(0) fib_notify(0) seq(0) lock(3)
```

Glossary

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

Key Term or Acronym	Description
BGP-Free Core	Network design where MPLS transport is used in the core without running BGP between core routers. Edge devices handle route advertisement and resolution.
Routing Information Base (RIB)	The routing table in control plane that maintains best route selections and recursive next-hop information.
Forwarding Information Base (FIB)	The hardware forwarding table that programs routes installed by the RIB for packet forwarding.
Recursive Resolution	The process of resolving a route's next-hop through another route or transport (e.g., MPLS tunnel) instead of a directly connected IP interface.
Multiprotocol Label Switching (MPLS)	Label-based forwarding technology used for scalable and efficient transport within the core network.
Next Hop Label Forwarding Entry (NHLFE)	MPLS data structure that defines label operations (push, swap, pop) and the next-hop interface for labeled forwarding.
FEC-to-NHLFE Mapping (FTN)	Mapping between Forwarding Equivalence Class (FEC) and NHLFE entries used to associate IP prefixes with MPLS label paths.
Label Distribution Protocol (LDP)	Protocol that distributes labels for MPLS forwarding paths dynamically across routers.
Resource Reservation Protocol - Traffic Engineering (RSVP-TE)	Protocol used to establish MPLS-TE LSPs with defined resource and path constraints.
Segment Routing - MPLS (SR-MPLS)	MPLS forwarding model where segment identifiers (SIDs) define packet paths without per-flow signaling.
Equal-Cost Multi-Path (ECMP)	Mechanism that enables load sharing across multiple equal-cost routes.
Fast Reroute (FRR)	Protection mechanism that provides rapid failover by pre-installing backup label paths.
Topology Independent Loop-	Advanced FRR mechanism ensuring near-instantaneous recovery for link or

Key Term or Acronym	Description
Free Alternate (TI-LFA)	node failures.
Hardware Support Layer (HSL)	OcNOS subsystem responsible for programming and managing hardware forwarding entries in FIB.
FTN-ID	Identifier assigned to an MPLS Forwarding Equivalence Class entry used in recursive resolution.
NHLFE-IX	Index that uniquely identifies the NHLFE used by the MPLS forwarding path.
BGP Multipath	Capability that allows installation of multiple equal-cost BGP paths; not supported in current release for BGP-Free Core scenarios.

© 2026 IP Infusion Inc. Proprietary

BORDER GATEWAY PROTOCOL COMMAND REFERENCE

BGP Commands 1244
address-family ipv4 1248
aggregate-address 1250
auto-summary 1252
bgp aggregate-nexthop-check 1253
bgp always-compare-med 1254
bgp as-local-count 1255
bgp bestpath as-path ignore 1256
bgp bestpath as-path multipath-relax 1257
bgp bestpath compare-confed-aspath 1258
bgp bestpath compare-routerid 1259
bgp bestpath dont-compare-originator-id 1260
bgp bestpath med 1261
bgp bestpath tie-break-on-age 1263
bgp client-to-client reflection 1264
bgp cluster-id 1265
bgp confederation vrf identifier 1266
bgp confederation vrf peers 1267
bgp dampening 1268
bgp default local-preference 1270
bgp deterministic-med 1271
bgp enforce-first-as 1273
bgp extended-asn-cap 1274
bgp fast-external-failover 1275
bgp implicit-null 1276
bgp log-neighbor-changes 1278
mpls-nexthop-tracking 1280
bgp nexthop-trigger delay 1281
bgp nexthop-trigger enable 1282
bgp rfc1771-path-select 1283
bgp rfc1771-strict 1284
bgp router-id 1285
bgp scan-time 1286
table-map 1287
bgp update-delay 1289

© 2026 IP Infusion Inc. Proprietary

bgp unnumbered-mode 1290
clear bgp (A.B.C.D|X:X::X:X) 1291
clear bgp * 1292
clear bgp * l2vpn evpn 1294
clear bgp <1-4294967295> 1295
clear bgp <1-4294967295> l2vpn evpn 1297
clear bgp A.B.C.D l2vpn evpn 1298
clear bgp dampening 1299
clear bgp external 1300
clear bgp flap-statistics 1302
clear bgp peer-group 1303
clear bgp peer-group WORD l2vpn evpn 1305
clear bgp statistics 1306
clear ip bgp * (description LINE|)* 1307
clear ip bgp (A.B.C.D|X:X::X:X|WORD) (description LINE|) 1308
clear ip bgp A.B.C.D 1309
clear ip bgp A.B.C.D vrf 1311
clear ip bgp all vrf (VRFNAME|all|default) (description LINE|) 1312
clear ip bgp table-map 1313
clear ip bgp vrf WORD (A.B.C.D|X:X::X:X|WORD) (description LINE|) 1314
debug bgp 1315
distance 1317
exit-address-family 1319
ip as-path access-list 1320
ip community-list <1-99> 1321
ip community-list <100-500> 1322
ip community-list expanded 1323
ip community-list standard 1324
ip community-list WORD 1326
ip extcommunity-list <1-99> 1328
ip extcommunity-list <100-500> 1329
ip extcommunity-list expanded 1330
ip extcommunity-list standard 1331
l2vpn-unnumbered-mode 1333
match ip peer 1334
max-paths 1335
max-paths eigbp 1336
neighbor activate 1338
neighbor advertisement-interval 1339
neighbor allowas-in 1340

© 2026 IP Infusion Inc. Proprietary

neighbor as-origination-interval 1342
neighbor attribute-unchanged 1343
neighbor authentication-key 1345
neighbor capability graceful-restart 1347
neighbor capability orf prefix-list 1348
neighbor collide-established 1350
neighbor de-activate 1352
neighbor default-originate 1354
neighbor description 1356
neighbor disallow-infinite-holdtime 1358
neighbor distribute-list 1359
neighbor ebgp-multihop 1361
neighbor enforce-multihop 1363
neighbor extended-optional-param 1364
neighbor fall-over bfd 1365
neighbor filter-list 1366
neighbor g-shut 1368
neighbor g-shut-timer 1369
neighbor limit 1370
neighbor local-as 1371
neighbor maximum-prefix 1373
neighbor next-hop-self 1375
neighbor optional-as 1376
neighbor override-capability 1377
neighbor passive 1378
neighbor peer-group 1379
neighbor WORD peer-group range 1383
neighbor port 1384
neighbor prefix-list 1385
neighbor remote-as 1387
neighbor remove-private-AS 1389
neighbor restart-time 1390
neighbor route-map 1391
neighbor route-reflector-client 1393
neighbor route-server-client 1395
neighbor send-community 1396
neighbor send-community large 1398
neighbor shutdown 1399
neighbor strict-capability-match 1401
neighbor timers 1402

© 2026 IP Infusion Inc. Proprietary

neighbor unsuppress-map 1404
neighbor update-source 1406
neighbor version 1407
neighbor weight 1409
neighbor WORD peer-group 1411
network 1412
network synchronization 1414
redistribute 1415
restart bgp graceful 1417
router bgp 1418
snmp restart bgp 1419
synchronization 1420
timers bgp 1421
v4-unnumbered-mode 1422
BGP4+ Commands 1423
address-family ipv6 1424
aggregate-address X:X::X:X/M 1426
clear bgp * ipv6 1428
clear bgp ipv6 (A.B.C.D|X:X::X:X) 1430
clear bgp ipv6 <1-4294967295> 1431
clear bgp ipv6 external 1432
clear bgp ipv6 peer-group 1433
clear bgp ipv6 unicast flap-statistics 1434
clear ipv6 bgp * vrf 1435
clear ipv6 bgp X:X::X:X vrf 1436
clear ip bpg ipv6 unicast table-map 1437
network X:X::X:X 1438
BGP Flowspec Commands 1440
address-family 1441
address-family ipv4 flowspec vrf 1442
class-map type traffic match-all 1443
drop 1444
flowspec 1445
hardware filter ipv4-bgp-flowspec 1446
ipv4 flowspec-disable 1447
local-install 1451
match destination-address 1452
match destination-port 1453
match dscp 1454
match fragment-type 1455

© 2026 IP Infusion Inc. Proprietary

match icmp-code 1456
match icmp-type 1457
match packet-length 1458
match port 1459
match protocol 1460
match source-address 1461
match source-port 1462
match tcp-flag 1463
neighbor flowspec-vaildation-disable 1464
police rate 1465
policy-map 1466
policy-map class type traffic 1467
redirect 1468
service-policy type pbr 1469
set dscp 1470
set ip next-hop 1471
set ipv6 next-hop 1472
show hsl bgp-flowspec-rules fib-id 1473
show ip bgp flowspec 1475
vrf 1481
BGP Virtual Private Network Commands 1482
bgp external-route-leak 1483
bgp inbound-route-filter 1484
clear ip bgp <1-4294967295> vpnv4 1485
clear ip bgp A.B.C.D vpnv4 1486
clear bgp <1-4294967295> l2vpn vpls 1487
clear bgp A.B.C.D l2vpn vpls 1488
clear bgp * l2vpn vpls 1489
clear ip bgp * vpnv4 1490
debug bgp mpls 1491
export map 1492
import map 1495
ip vrf 1496
neighbor allow-ebgp-vpn 1497
neighbor as-override 1498
neighbor send-community 1499
neighbor soo 1500
rd (route distinguisher) 1501
route-target 1502
BGP Show Commands 1504

© 2026 IP Infusion Inc. Proprietary

show bgp 1506
show bgp A.B.C.D 1507
show bgp A.B.C.D/M 1509
show bgp client 1511
show bgp community 1512
show bgp community-list 1514
show bgp dampening dampened-paths 1515
show bgp dampening flap-statistics 1517
show bgp dampening parameters 1519
show bgp filter-list 1521
show bgp inconsistent-as 1523
show bgp ipv6 1524
show bgp ipv6 peer-group 1527
show bgp l2vpn vpls 1528
show bgp label-pool 1531
show bgp neighbors 1532
show bgp neighbors advertised-routes 1536
show bgp neighbors received prefix-filter 1537
show bgp neighbors received-routes 1538
show bgp neighbors routes 1540
show bgp nexthop-tracking 1542
show bgp nexthop-tree-details 1544
show bgp paths 1545
show bgp prefix-list 1546
show bgp quote-regexp 1548
show bgp regexp 1549
show bgp route-map 1550
show bgp statistics 1552
show bgp summary 1554
show bgp X:X::X:X 1557
show bgp X:X::X:X/M longer prefixes 1558
show debugging bgp 1559
show ip bgp 1560
show ip bgp cidr-only 1563
show ip bgp community-info 1564
show ip bgp peer-group 1565
show ip bgp peer-group vrf all 1566
show ip bgp rtfilter all 1567
show ip bgp rtfilter neighbors 1570
show ip bgp scan 1572

© 2026 IP Infusion Inc. Proprietary

show ip bgp vpnv4 1573
show ip bgp vpnv6 all neighbors 1579
show ip bgp vpnv6 rd neighbors 1582
show ip extcommunity-list 1584
show ip protocols 1585
show ip vrf 1586
show running-config as-path access-list 1587
show running-config bgp 1588
show running-config community-list 1589
show running-config vrf 1590
show running-config ip vrf 1591
show running-config mac vrf 1592
show running-config vrf word 1593
BGP Additional Paths Commands 1594
bgp additional-paths 1595
bgp additional-paths install 1596
neighbor additional-paths 1597
neighbor advertise additional-paths 1598
BGP Graceful Restart Commands 1599
bgp g-shut 1600
bgp g-shut-capable 1601
bgp g-shut-local-preference 1602
bgp graceful-restart 1603
bgp update-delay 1605
neighbor capability graceful-restart 1606
neighbor g-shut 1608
neighbor g-shut-timer 1609
neighbor restart-time 1610
BGP Labeled Unicast Commands 1611
address-family labeled-unicast 1612
allocate-label 1614
clear ip bgp * ipv4 labeled-unicast 1615
clear ip bgp A.B.C.D ipv4 labeled-unicast 1616
clear ip bgp peer-group WORD ipv4 labeled-unicast 1617
Regular Expressions 1618

© 2026 IP Infusion Inc. Proprietary

BGP Commands

This section provides an organized overview of CLI commands used to configure and manage Border Gateway Protocol (BGP), including both BGP4 for IPv4 and BGP4+ for IPv6 and other network layer protocols.

Refer to BGP4+ Commands (page 1423) section for IPv6 commands.

BGP/BGP4 commands apply to IPv4 routing and are placed directly under the BGP routing process. BGP4 is used to exchange routing information between Autonomous Systems (AS) in IPv4-based networks. It includes the following commands:

address-family ipv4 1248
aggregate-address 1250
auto-summary 1252
bgp aggregate-nexthop-check 1253
bgp always-compare-med 1254
bgp as-local-count 1255
bgp bestpath as-path ignore 1256
bgp bestpath as-path multipath-relax 1257
bgp bestpath compare-confed-aspath 1258
bgp bestpath compare-routerid 1259
bgp bestpath dont-compare-originator-id 1260
bgp bestpath med 1261
bgp bestpath tie-break-on-age 1263
bgp client-to-client reflection 1264
bgp cluster-id 1265
bgp confederation vrf identifier 1266
bgp confederation vrf peers 1267
bgp dampening 1268
bgp default local-preference 1270
bgp deterministic-med 1271
bgp enforce-first-as 1273
bgp extended-asn-cap 1274
bgp fast-external-failover 1275
bgp implicit-null 1276
bgp log-neighbor-changes 1278
mpls-nexthop-tracking 1280
bgp nexthop-trigger delay 1281
bgp nexthop-trigger enable 1282
bgp rfc1771-path-select 1283
bgp rfc1771-strict 1284
bgp router-id 1285

© 2026 IP Infusion Inc. Proprietary

bgp scan-time 1286
table-map 1287
bgp update-delay 1289
bgp unnumbered-mode 1290
clear bgp (A.B.C.D|X:X::X:X) 1291
clear bgp * 1292
clear bgp * l2vpn evpn 1294
clear bgp <1-4294967295> 1295
clear bgp <1-4294967295> l2vpn evpn 1297
clear bgp A.B.C.D l2vpn evpn 1298
clear bgp dampening 1299
clear bgp external 1300
clear bgp flap-statistics 1302
clear bgp peer-group 1303
clear bgp peer-group WORD l2vpn evpn 1305
clear bgp statistics 1306
clear ip bgp * (description LINE|)* 1307
clear ip bgp (A.B.C.D|X:X::X:X|WORD) (description LINE|) 1308
clear ip bgp A.B.C.D 1309
clear ip bgp A.B.C.D vrf 1311
clear ip bgp all vrf (VRFNAME|all|default) (description LINE|) 1312
clear ip bgp table-map 1313
clear ip bgp vrf WORD (A.B.C.D|X:X::X:X|WORD) (description LINE|) 1314
debug bgp 1315
distance 1317
exit-address-family 1319
ip as-path access-list 1320
ip community-list <1-99> 1321
ip community-list <100-500> 1322
ip community-list expanded 1323
ip community-list standard 1324
ip community-list WORD 1326
ip extcommunity-list <1-99> 1328
ip extcommunity-list <100-500> 1329
ip extcommunity-list expanded 1330
ip extcommunity-list standard 1331
l2vpn-unnumbered-mode 1333
match ip peer 1334
max-paths 1335
max-paths eigbp 1336

© 2026 IP Infusion Inc. Proprietary

neighbor activate 1338
neighbor advertisement-interval 1339
neighbor allowas-in 1340
neighbor as-origination-interval 1342
neighbor attribute-unchanged 1343
neighbor authentication-key 1345
neighbor capability graceful-restart 1347
neighbor capability orf prefix-list 1348
neighbor collide-established 1350
neighbor de-activate 1352
neighbor default-originate 1354
neighbor description 1356
neighbor disallow-infinite-holdtime 1358
neighbor distribute-list 1359
neighbor ebgp-multihop 1361
neighbor enforce-multihop 1363
neighbor extended-optional-param 1364
neighbor fall-over bfd 1365
neighbor filter-list 1366
neighbor g-shut 1368
neighbor g-shut-timer 1369
neighbor limit 1370
neighbor local-as 1371
neighbor maximum-prefix 1373
neighbor next-hop-self 1375
neighbor optional-as 1376
neighbor override-capability 1377
neighbor passive 1378
neighbor peer-group 1379
neighbor WORD peer-group range 1383
neighbor port 1384
neighbor prefix-list 1385
neighbor remote-as 1387
neighbor remove-private-AS 1389
neighbor restart-time 1390
neighbor route-map 1391
neighbor route-reflector-client 1393
neighbor route-server-client 1395
neighbor send-community 1396
neighbor send-community large 1398

© 2026 IP Infusion Inc. Proprietary

neighbor shutdown 1399
neighbor strict-capability-match 1401
neighbor timers 1402
neighbor unsuppress-map 1404
neighbor update-source 1406
neighbor version 1407
neighbor weight 1409
neighbor WORD peer-group 1411
network 1412
network synchronization 1414
redistribute 1415
restart bgp graceful 1417
router bgp 1418
snmp restart bgp 1419
synchronization 1420
timers bgp 1421
v4-unnumbered-mode 1422

© 2026 IP Infusion Inc. Proprietary

address-family ipv4

Use the address family command to enter the IPv4 or VPNv4 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 l2vpn evpn
address-family rtfilter unicast
address-family vpn4
address-family vpn4 unicast
no address-family ipv4 vrf NAME
no address-family ipv4 (unicast|multicast)
no address-family l2vpn evpn
no address-family rtfilter unicast
no address-family vpnv4
no address-family vpnv4 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

l2vpn evpn
Layer 2 VPN routing sessions with EVPN endpoint information distributed to BGP peers

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.

© 2026 IP Infusion Inc. Proprietary

unicast
Unicast address prefixes

vpnv4
VPN version 4 address family

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

```
(config)#router bgp 100
(config-router)#address-family ipv4 unicast
(config-router-af)#no aggregate-address 10.0.0.0/8
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1256) 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1285) 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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 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.

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp confederation vrf vrf1 identifier 1000
```

© 2026 IP Infusion Inc. Proprietary

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.

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 the below 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
```

© 2026 IP Infusion Inc. Proprietary

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

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

bgp implicit-null

Use this command to assign the MPLS label implicit-null (with a label value of 3) for locally originated IPv4 routes, whether they are connected or redistributed from other protocols which are found in the BGP IPv4 Unicast RIB.

Use the `no` form of this command to disable the default implicit-null behavior by configuring `no bgp implicit-null` under `address-family ipv4 labeled-unicast.` The `no bgp implicit-null` CLI command prompts BGP to assign a valid label (with a value greater than 16) to all IPv4 prefixes.

Command Syntax

```
bgp implicit-null
no bgp implicit-null
```

Parameters

None

Default

The `bgp implicit-null` is enabled under `address-family ipv4 labeled-unicast`.

Command Mode

BGP address-family ipv4 labelled-unicast mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

```
OcNOS#configure terminal
```

Note: Describe the show output fields for new show commands. (Example: Refer the command `show streaming-telemetry dynamic-subscriptions`).

To allocate an MPLS implicit-null label to IPv4 unicast prefix 22.2.2.2/32s

```
!
router bgp 100
bgp router-id 2.2.2.2
allocate-label all
neighbor 1.1.1.1 remote-as 100
neighbor 1.1.1.1 update-source lo
!
address-family ipv4 unicast
network 22.2.2.2/32
exit-address-family
!
address-family ipv4 labeled-unicast
bgp implicit-null
neighbor 1.1.1.1 activate
```

© 2026 IP Infusion Inc. Proprietary

```
exit-address-family
!
exit
!
```

Example for non-implicit-null

To allocate a valid MPLS unique label (> 16) to IPv4 unicast prefix 22.2.2.2/32

```
!
router bgp 100
bgp router-id 2.2.2.2
allocate-label all
neighbor 1.1.1.1 remote-as 100
neighbor 1.1.1.1 update-source lo
!
address-family ipv4 unicast
network 22.2.2.2/32
exit-address-family
!
address-family ipv4 labeled-unicast
no bgp implicit-null
neighbor 1.1.1.1 activate
exit-address-family
!
exit
!
```

© 2026 IP Infusion Inc. Proprietary

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

© 2026 IP Infusion Inc. Proprietary

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
(config)#router bgp 100
(config-router)#bgp log-neighbor-changes
```

© 2026 IP Infusion Inc. Proprietary

mpls-nexthop-tracking

Use this command to register and track nexthops of BGP prefixes for MPLS reachability resolution.

Use “no” form of this command to deactivate MPLS nexthop tracking for BGP routes.

Note: `bgp nexthop-trigger enable` must be configured to enable `mpls-nexthop-tracking`

Command Syntax

```
mpls-nexthop-tracking
no mpls-nexthop-tracking
```

Parameters

None

Command Mode

BGP mode

Applicability

This command is introduced in OcNOS version 5.0 and is updated in OcNOS version 6.4.0

MPLS NHT can be used to track nexthop LSP's of BGP prefixes, MPLS NHT supports BGP IPv4 Labeled Unicast, VPNv4, VPNv6 (6VPE), 6PE routes.

Examples

```
#configure terminal
(config)#bgp nexthop-trigger enable
(config)#router bgp 100
(config-router)#mpls-nexthop-tracking
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 modeConfigure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)# bgp nexthop-trigger enable
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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:

When an explicit BGP router-id is configured/un-configured

When the router's (global) router-id is set/unset (holds true when (a) is not applicable),

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

table-map

Use this command to enable or disable suppression/modification of incoming BGP updates to IP RIB/FIB table installation.

Note: The `table-map filter` configuration in the VRF address family applies only to RIB routes and not to MPLS routes.

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 version 4.2.

Examples

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#
```

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 execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear bgp 3.3.3.3
```

© 2026 IP Infusion Inc. Proprietary

clear bgp *

Use this command to reset the BGP connection for all peers.

Note: The auto-clear soft feature that supports route-map modifications, such as new match, set add, replace, or delete does not support the updation of internally associated MAC-lists. For this, the admin must use the explicit `clear bgp * soft in/out` command.

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.

© 2026 IP Infusion Inc. Proprietary

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 execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear bgp *
#clear ip bgp * ipv4 unicast in prefix-filter
```

© 2026 IP Infusion Inc. Proprietary

clear bgp * l2vpn evpn

Use this command to reset the BGP L2VPN EVPN connection for all the BGP peers.

Command Syntax

```
clear bgp * l2vpn evpn soft (in | out | )
```

Parameters

soft
Clear both incoming and outgoing routes.

in
Soft reconfig inbound update.

out
Soft reconfig outbound update.

Default

None.

Command Mode

Privileged execution mode

Applicability

This command was introduced in OcNOS version 6.3.4.

Examples

```
OcNOS#clear bgp * l2vpn evpn soft
OcNOS#clear bgp * l2vpn evpn soft in
```

© 2026 IP Infusion Inc. Proprietary

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.

in
Clear incoming advertised routes.

prefix-filter
Push out prefix-list ORF and do inbound soft reconfig.

© 2026 IP Infusion Inc. Proprietary

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 execution 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 execution 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
```

© 2026 IP Infusion Inc. Proprietary

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 execution 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
```

© 2026 IP Infusion Inc. Proprietary

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 execution 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
```

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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 execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear ip bgp external
```

© 2026 IP Infusion Inc. Proprietary

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 execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip bgp flap-statistics
#clear ip bgp ipv4 unicast flap-statistics
```

© 2026 IP Infusion Inc. Proprietary

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.

in
Clear incoming advertised routes.

© 2026 IP Infusion Inc. Proprietary

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 execution 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 execution 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
```

© 2026 IP Infusion Inc. Proprietary

clear bgp statistics

Use this command to reset all BGP statistics.

Command Syntax

```
clear bgp statistics
```

Parameters

None

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear bgp statistics
```

© 2026 IP Infusion Inc. Proprietary

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 execution 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 execution 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
```

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

soft
Clear both incoming and outgoing routes.

in
Soft reconfig inbound update.

out
Soft reconfig outbound update.

Command Mode

Privileged execution 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 execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip bgp 35.0.0.1 vrf
```

© 2026 IP Infusion Inc. Proprietary

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 execution 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
```

© 2026 IP Infusion Inc. Proprietary

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 execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip bgp table-map vrf all
```

© 2026 IP Infusion Inc. Proprietary

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 execution 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
```

© 2026 IP Infusion Inc. Proprietary

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 mpls
debug bgp nht
debug bgp nsm
debug bgp updates
debug bgp updates (in|out)
debug bgp vpls
debug bgp evpn
debug bgp route-process
no debug bgp evpn
no debug bgp route-process
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 mpls
no debug bgp nht
no debug bgp nsm
no debug bgp updates
no debug bgp vpls
```

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

© 2026 IP Infusion Inc. Proprietary

mpls
Enable debugging for BGP Multiprotocol Label Switching (MPLS)

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

vpls
Enable debugging for BGP Virtual Private LAN Service (VPLS)

Command Mode

Privileged execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3. Introduced the following options:

- debug bgp evpn
- debug bgp route-process
- no debug bgp evpn
- no debug bgp route-process

and removed undebug options in the release OcNOS version 6.6.0

Examples

```
#debug bgp
#debug bgp events
```

© 2026 IP Infusion Inc. Proprietary

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 {intra-area |inter-area |external}
```

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.

© 2026 IP Infusion Inc. Proprietary

Example

```
#configure terminal
(config)#router ipv6 ospf 100
(config-router)#distance ospfv3 inter-area 20 intra-area 10 external 40
```

© 2026 IP Infusion Inc. Proprietary

exit-address-family

Use this command to exit Address-Family-vrf, Address-Family-vpnv4 mode.

For information on how to enter the address family mode (IPv4, VPNv4), see .

Command Syntax

```
exit-address-family
```

Parameters

None

Default

No default value is specified

Command Mode

Address Family-vrf and Address Family-vpnv4 mode.

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)#
```

© 2026 IP Infusion Inc. Proprietary

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 (deny | permit) LINE
```

Parameters

WORD
Access list name

deny
Reject packets

permit
Forward packets

LINE
An ordered list as a regular expression.

Note: In regular expression, for number value use {0,n} instead of {,n}.

Default

No default value is specified

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3 and updated in OcNOS version 6.4.

Examples

```
#configure terminal
(config)#ip as-path access-list mylist deny ^65535$
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 .*
```

© 2026 IP Infusion Inc. Proprietary

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 .*
```

© 2026 IP Infusion Inc. Proprietary

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

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

Examples

```
#configure terminal
(config)#ip community-list mylist permit 7675:80 7675:90
(config)#no ip community-list mylist permit 7675:80 7675:90
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 *.
```

© 2026 IP Infusion Inc. Proprietary

ip extcommunity-list expanded

Use this command to create an expanded community list. The expanded community-list defines the community attributes with regular expressions. The expanded community-list match will happen based on input given for regular expressions.

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 e1 permit 2001:1 300:1
(config)#ip extcommunity-list expanded CLIST permit ^[2001:1 300:1]$
```

© 2026 IP Infusion Inc. Proprietary

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)
no ip extcommunity-list standard WORD 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
```

© 2026 IP Infusion Inc. Proprietary

```
(config)#ip extcommunity-list standard 36 permit rt 5675:50
(config)#ip extcommunity-list standard CLIST permit soo 10.10.32.15:70
```

© 2026 IP Infusion Inc. Proprietary

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 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)#
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

```
(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
```

© 2026 IP Infusion Inc. Proprietary

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 and VPNv4 address prefix types, 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 (page 1411) 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)#address-family vpnv4 unicast
(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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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, VPNv4 unicast, VPNv6 unicast, IPv4/IPv6 VRF, L2VPN EVPN`), with the following limitation:

– Under `address-family ipv4 labeled-unicast, neighbor allowas-in` is supported for individual neighbors only; it is not supported for `peer-group neighbors` (neighbor <peer-group-name>). Use per-neighbor configuration instead of `neighbor <peer-group-name> allowas-in.`

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

```
(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
```

© 2026 IP Infusion Inc. Proprietary

```
(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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1603) 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 (page 1411) 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 - Applicable for ipv4, ipv6 and vpnv4

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

(config-router)#address-family vpnv4 unicast
(config-router-af)#neighbor 2.2.2.2 activate
(config-router-af)#neighbor 2.2.2.2 capability graceful-restart
(config-router-af)#neighbor 3.3.3.3 activate
(config-router-af)#neighbor 3.3.3.3 capability graceful-restart
(config-router-af)#exit-address-family
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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, ipv4 labeled-unicast, vpnv4 unicast, v4-unnumbered, ipv6 unicast, ipv6 labeled-unicast, vpnv6 unicast) mode

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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.

© 2026 IP Infusion Inc. Proprietary

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 collide-established
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

```
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
!
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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.

© 2026 IP Infusion Inc. Proprietary

```
#configure 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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.

© 2026 IP Infusion Inc. Proprietary

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 description unnbered-peer1
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 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 (page 1411) 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.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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.

© 2026 IP Infusion Inc. Proprietary

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 ebgp-multihop
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 <1-512>
no neighbor WORD (limit <1-512>|)
```

Parameters

WORD

Name of a BGP peer group created with the neighbor WORD peer-group (page 1411) command.

<1-512>

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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 (page 1309)` 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.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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.

Notes:

- See Usage (page 1380) below for when a peer group and a neighbor have conflicting attribute configurations.
- When a peer is detached from a peer-group and if an attempt is made to create a static peer with the same address immediately within the single commit transaction, then the system throws the error: `%% Peer deletion is in progress, peer cannot be created`. For example, refer to the Peer Creation Error During Deletion (page 1380) section.

To create a peer group, use the `neighbor WORD peer-group (page 1411)` 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.

© 2026 IP Infusion Inc. Proprietary

Example

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor group1 peer-group
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

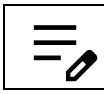

Note: For unnumbered peer we use peergroup key word to bind the unnumbered peer in group.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 peergroup ipi
```

Peer Creation Error During Deletion

Attempting to re-create a BGP peer immediately after removing it from a peer-group within the same commit:

```
OcNOS#show running-config router bgp
!
router bgp 65000
 neighbor RR peer-group
 neighbor RR remote-as 65000
 neighbor RR update-source loopback200
 neighbor 223.223.223.2 peer-group RR
 neighbor 223.223.223.3 peer-group RR
 !
 address-family vpnv4 unicast
 neighbor RR activate
 exit-address-family
 !
 exit
!

OcNOS#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
OcNOS(config)#router bgp 65000
OcNOS(config-router)#no neighbor 223.223.223.2 peer-group RR
OcNOS(config-router)#neighbor 223.223.223.2 remote-as 65000
OcNOS(config-router)#neighbor 223.223.223.2 update-source loopback200
OcNOS(config-router)#address-family vpnv4 unicast
OcNOS(config-router-af)#neighbor 223.223.223.2 activate
OcNOS(config-router-af)#commit
%% Peer deletion is in progress, peer cannot be created -  /bgp/bgp-instances/bgp-instance[bgp-
as='65000']/peers/peer[peer-address='223.223.223.2']/config/peer-as
%% bgpd config validation failed
Error
 %% Delete not supported at the Peer Address-family. Delete at the global address-family
% Failed to commit .. As error(s) encountered during commit operation...
  Uncommitted configurations are retained in the current transaction session, check 'show transaction
current'.
  Correct the reason for the failure and re-issue the commit.
  Use 'abort transaction' to terminate current transaction session and discard all uncommitted
changes.
%% Peer deletion is in progress, peer cannot be created
OcNOS(config-router-af)#
```

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 (page 1338)
- neighbor advertisement-interval (page 1339)
- neighbor as-origination-interval (page 1342)
- neighbor attribute-unchanged (page 1343)
- neighbor capability orf prefix-list (page 1348)
- neighbor distribute-list (page 1359) with an access-list number and the `out` parameter
- neighbor dont-capability-negotiate
- neighbor filter-list (page 1366) with the `out` parameter
- neighbor next-hop-self (page 1375)
- neighbor prefix-list (page 1385) with an access-list name and the `out` parameter
- neighbor remove-private-AS (page 1389)
- neighbor route-map (page 1391) with the `out` parameter
- neighbor route-reflector-client (page 1393)
- neighbor route-server-client (page 1395)
- neighbor send-community (page 1396)
- neighbor unsuppress-map (page 1404)

These commands control inbound attribute updates:

- neighbor allowas-in (page 1340)
- neighbor collide-established (page 1350)
- neighbor description (page 1356)
- neighbor distribute-list (page 1359) with an access-list number and the `in` parameter
- neighbor ebgp-multihop (page 1361)
- neighbor enforce-multihop (page 1363)
- neighbor filter-list (page 1366) with the `in` parameter
- neighbor g-shut (page 1368)
- neighbor g-shut-timer (page 1369)
- neighbor local-as (page 1371)
- neighbor maximum-prefix (page 1373)
- neighbor override-capability (page 1377)
- neighbor passive (page 1378)
- neighbor authentication-key (page 1345)
- neighbor port (page 1384)

© 2026 IP Infusion Inc. Proprietary

- neighbor prefix-list (page 1385) with an access-list name and the `in` parameter
- neighbor remote-as (page 1387)
- neighbor restart-time (page 1390)
- neighbor route-map (page 1391) with the `in` parameter
- neighbor shutdown (page 1399)
- neighbor strict-capability-match (page 1401)
- neighbor update-source (page 1406)
- neighbor weight (page 1409)

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.
- Users cannot bind a peer to a peer group if it is already activated in address-family mode. Therefore, it is necessary to deactivate the peer first.

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1359) 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 (page 1411) 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.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1338) 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:

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 (page 1411) 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, 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.

© 2026 IP Infusion Inc. Proprietary

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
    !
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 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 (page 1603) 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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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.

Note: Route maps enable filtering and modification of BGP routes using 'match' and 'set' clauses. They can be applied to a BGP peer or peer-group in either the inbound or outbound direction to influence received or advertised routes. However, for route-map updates to take effect, a manual command such as clear ip bgp A.B.C.D (page 1309) `<neighbor> soft` must be executed for each BGP peer or peer-group address-family. This manual step is required in all software releases except for release 6.6.x and later, provided that the 'bgp auto-policy-soft-reset enable' feature is configured.

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 (page 1411) 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

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1268) 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 (page 1499) in 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 (page 1411) 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.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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 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 “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
```

© 2026 IP Infusion Inc. Proprietary

```
(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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

neighbor unsuppress-map

Use this command to selectively leak more-specific routes to a particular neighbor.

When the aggregate-address (page 1250) 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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

```
(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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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.

© 2026 IP Infusion Inc. Proprietary

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 version 4
```

© 2026 IP Infusion Inc. Proprietary

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/VPNv4/6PE 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 (page 1411) 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.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

None

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

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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|connected-host-routes|ospf|rip|static]
redistribute [connected|isis|connected-host-routes|ospf|rip|static] route-map WORD
no redistribute [connected|isis|connected-host-routes|ospf|rip|static]
no redistribute [connected|isis|connected-host-routes|ospf|rip|static] route-map WORD
```

Parameters

connected
Redistribute connected routes

isis
Redistribute connected ISO IS-IS routes

connected-host-routes
Redistribute connected host to routes

ospf
Redistribute OSPFv2 routes

rip
Redistribute RIP routes

static
Redistribute static routes

route-map
Route map reference

WORD
Route map entires

Default

Disabled

Command Mode

Address Family mode and Address Family-vrf mode

Applicability

This command was introduced before OcNOS version 1.3.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

Disabled

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#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
```

© 2026 IP Infusion Inc. Proprietary

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)#
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 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)#
```

© 2026 IP Infusion Inc. Proprietary

BGP4+ Commands

BGP4+ extends BGP Commands (page 1244) to support IPv6 and other network layer protocols using address-family configuration mode. IPv6-specific commands are encapsulated within the address-family ipv6 context. It includes the following commands:

address-family ipv6 1424
aggregate-address X:X::X:X/M 1426
clear bgp * ipv6 1428
clear bgp ipv6 (A.B.C.D|X:X::X:X) 1430
clear bgp ipv6 <1-4294967295> 1431
clear bgp ipv6 external 1432
clear bgp ipv6 peer-group 1433
clear bgp ipv6 unicast flap-statistics 1434
clear ipv6 bgp * vrf 1435
clear ipv6 bgp X:X::X:X vrf 1436
clear ip bpg ipv6 unicast table-map 1437
network X:X::X:X 1438

© 2026 IP Infusion Inc. Proprietary

address-family ipv6

Use the address family command to enter the IPv6 or VPNv6 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-IPv6 and IPv6 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-IPv6 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 IPv6 routes.

Use the `no` parameter with this command to disable the address-family configurations.

Command Syntax

```
address-family ipv6 (unicast|)
address-family ipv6 vrf NAME
address-family vpn6
address-family vpn6 unicast
no address-family ipv6 vrf NAME
no address-family ipv6 (unicast|)
no address-family vpnv6
no address-family vpnv6 unicast
```

Parameters

ipv6
Specifies the IPv6 address family.

unicast
Specifies the unicast address prefixes.

vrf NAME
Specifies the Virtual Private Network (VPN) routing/forwarding instance name.

vpnv6
Specifies the VPN version 6 (v6) address family.

vpnv6 unicast
Specifies the VPNv6 unicast address prefixes.

Applicability

None

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

© 2026 IP Infusion Inc. Proprietary

Examples

```
#configure terminal
(config)#router bgp 7657
(config-router)#address-family ipv6 unicast
(config-router-af)#exit-address-family

(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
(config-router-af)#exit-address-family
```

© 2026 IP Infusion Inc. Proprietary

aggregate-address X:X::X:X/M

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. This 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 more-specific routes to neighbors.

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 the `as-set` 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 X:X::X:X/M
aggregate-address X:X::X:X/M as-set
aggregate-address X:X::X:X/M as-set summary-only
aggregate-address X:X::X:X/M summary-only
aggregate-address X:X::X:X/M summary-only as-set
no aggregate-address X:X::X:X/M
no aggregate-address X:X::X:X/M as-set
no aggregate-address X:X::X:X/M as-set summary-only
no aggregate-address X:X::X:X/M summary-only
no aggregate-address X:X::X:X/M summary-only as-set
```

Parameters

X:X::X:X/M
Aggregate IPv6 prefix

as-set
Generate AS set path information

summary-only
Filter more specific routes from updates

Command Mode

Address Family mode

Default

By default, aggregate address is disabled

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 100
```

© 2026 IP Infusion Inc. Proprietary

```
(config-router)#neighbor 2.2.2.2 remote-as 100
(config-router)#neighbor 3.3.3.3 remote-as 200
(config-router)#address-family ipv6
(config-router-af)#aggregate-address 3ffe::/32 as-set summary-only
```

clear bgp * ipv6

Use this command to reset the BGP IPv6 connection for all peers.

Note: The `soft in` or `in` and `soft out` or `out` in the BGP commands performs the same functionality. User can use any of the commands to soft reset.

Command Syntax

```
clear bgp ipv6 *
clear bgp ipv6 * in
clear bgp ipv6 * in prefix-filter
clear bgp ipv6 * out
clear bgp ipv6 * soft
clear bgp ipv6 * soft in
clear bgp ipv6 * soft out
clear ip bgp * ipv6 unicast in
clear ip bgp * ipv6 unicast out
clear ip bgp * ipv6 unicast soft
clear ip bgp * ipv6 unicast soft in
clear ip bgp * ipv6 unicast 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 all incoming and outgoing IPv6 routes.

in
Soft reconfig inbound update.

out
Soft reconfig outbound update.

unicast
Unicast prefixes

in
Clear incoming advertised routes.

out
Soft reconfig outbound update

soft
Clear both incoming and outgoing routes

in
Soft reconfig inbound update

out
Clear outgoing advertised routes

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip bgp * ipv6 unicast soft out
```

© 2026 IP Infusion Inc. Proprietary

clear bgp ipv6 (A.B.C.D|X:X::X:X)

Use this command to reset the BGP neighbor addresses (IPv4 or IPv6) for IPv6 peers.

Note: The `soft in` or `in` and `soft out` or `out` in the BGP commands performs the same functionality. User can use any of the commands to soft reset.

Command Syntax

```
clear bgp ipv6 (A.B.C.D|X:X::X:X)
clear bgp ipv6 (A.B.C.D|X:X::X:X) in
clear bgp ipv6 (A.B.C.D|X:X::X:X) in prefix-filter
clear bgp ipv6 (A.B.C.D|X:X::X:X) out
clear bgp ipv6 (A.B.C.D|X:X::X:X) soft
clear bgp ipv6 (A.B.C.D|X:X::X:X) soft in
clear bgp ipv6 (A.B.C.D|X:X::X:X) 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

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear bgp ipv6 10.5.2.7
```

© 2026 IP Infusion Inc. Proprietary

clear bgp ipv6 <1-4294967295>

Use this command to reset the BGP connection with a specified AS (Autonomous System) number for IPv6 peers.

Command Syntax

```
clear bgp ipv6 <1-4294967295>
clear bgp ipv6 <1-4294967295> in
clear bgp ipv6 <1-4294967295> in prefix-filter
clear bgp ipv6 <1-4294967295> out
clear bgp ipv6 <1-4294967295> soft
clear bgp ipv6 <1-4294967295> soft in
clear bgp ipv6 <1-4294967295> 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.

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear bgp ipv6 12345
```

© 2026 IP Infusion Inc. Proprietary

clear bgp ipv6 external

Use this command to reset the BGP connection for all external IPv6 peers or for a specified external IPv6 peer.

Command Syntax

```
clear bgp ipv6 external
clear bgp ipv6 external in prefix-filter
clear bgp ipv6 external soft
clear bgp ipv6 external soft in
clear bgp ipv6 external soft out
clear bgp ipv6 external WORD in
clear bgp ipv6 external WORD out
```

Parameters

in

Clear incoming advertised routes

prefix-filter

Push out prefix-list ORF and do inbound soft reconfig

soft

Clear both incoming and outgoing routes

in

Soft reconfig inbound update

out

Soft reconfig outbound update

WORD

Name of external IPv6 peer

in

Soft reconfig inbound update

out

Soft reconfig outbound update

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear bgp ipv6 external soft in
```

© 2026 IP Infusion Inc. Proprietary

clear bgp ipv6 peer-group

Use this command to reset the BGP connection for all members of a peer group.

Note: The `soft in` or `in` and `soft out` or `out` in the BGP commands performs the same functionality. User can use any of the commands to soft reset.

Command Syntax

```
clear bgp ipv6 peer-group WORD
clear bgp ipv6 peer-group WORD in
clear bgp ipv6 peer-group WORD in prefix-filter
clear bgp ipv6 peer-group WORD out
clear bgp ipv6 peer-group WORD soft
clear bgp ipv6 peer-group WORD soft in
clear bgp ipv6 peer-group WORD 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

WORD
BGP peer-group name

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip bgp peer-group P1
```

clear bgp ipv6 unicast flap-statistics

Use this command to reset IPv6 BGP route flap statistics.

Command Syntax

```
clear bgp ipv6 unicast flap-statistics
clear bgp ipv6 unicast flap-statistics X:X::X:X
clear bgp ipv6 unicast flap-statistics X:X::X:X/M
clear ip bgp ipv6 unicast flap-statistics
clear ip bgp ipv6 unicast flap-statistics X:X::X:X
clear ip bgp ipv6 unicast flap-statistics X:X::X:X/M
```

Parameters

X:X::X:X
IP prefix (network) for example, 35.0.0.0

X:X::X:X/M
IP prefix (network and length) for example, 35.0.0.0/8

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip bgp ipv6 unicast flap-statistics 1:2::3:4/7
```

© 2026 IP Infusion Inc. Proprietary

clear ipv6 bgp * vrf

Use this command to reset an IPv6 BGP connection for all virtual routing forwarding addresses.

Note: The `soft in` or `in` and `soft out` or `out` in the BGP commands performs the same functionality. User can use any of the commands to soft reset.

Command Syntax

```
clear ipv6 bgp * vrf WORD
clear ipv6 bgp * vrf WORD in
clear ipv6 bgp * vrf WORD out
clear ipv6 bgp * vrf WORD soft
clear ipv6 bgp * vrf WORD soft in
clear ipv6 bgp * vrf WORD soft out
```

Parameters

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

WORD
BGP peer-group name

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ipv6 bgp * vrf BGPpeer7
```

© 2026 IP Infusion Inc. Proprietary

clear ipv6 bgp X:X::X:X vrf

Use this command to reset the specified VPNv6 Routing/Forwarding (VRF) instance for BGP connections. If the neighbor address is specified with this command, it clears the specified connection. If no address is specified, this command clears all the BGP routes.

Note: The `soft in` or `in` and `soft out` or `out` in the BGP commands performs the same functionality. User can use any of the commands to soft reset.

Command Syntax

```
clear ipv6 bgp (X:X::X:X) vrf WORD
clear ipv6 bgp X:X::X:X vrf WORD in
clear ipv6 bgp X:X::X:X vrf WORD out
clear ipv6 bgp X:X::X:X vrf WORD soft in
clear ipv6 bgp X:X::X:X vrf WORD soft out
```

Parameters

vrf
Specify a VPN routing/forwarding instance for IPv6

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 execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ipv6 bgp vrf 10:10::0:12 vrf vrfname
```

clear ip bpg ipv6 unicast 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 bpg ipv6 unicast 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.

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip bgp ipv6 unicast table-map vrf all
```

network X:X::X:X

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` parameter with this command to remove an entry.

Command Syntax

```
network X:X::X:X/M
network X:X::X:X/M backdoor
network X:X::X:X/M route-map WORD (backdoor|)
no network X:X::X:X/M
no network X:X::X:X/M backdoor
no network X:X::X:X/M route-map WORD (backdoor|)
```

Parameters

X:X::X:X
IPv6 prefix <network>, for example, 3ffe::

backdoor
Specify a BGP backdoor route

WORD
Name of the route map

Default

None

Command Mode

IPv6 Address Family mode, IPv6 VRF Address Family mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#router bgp 10
(config-router)#network 172.26.0.0/16
```

If Router1 receives updates from 172.10.0.0 via two routing protocols RIP (distance 120) and eBGP (distance 20), router1 chooses the shorter route. Use the `backdoor` parameter to allow Router1 to learn about 172.10.0.0 via RIP.

```
(config)# router rip
(config)#network 172.10.0.0
(config)#router bgp 65001
(config)#neighbor 5ffe:507::2 remote-as 65001
(config-af)#address-family ipv6 unicast
(config-af)#network 8ffe:507::/64
(config-af)#neighbor 5ffe:507::2 activate
(config-af)#exit-address-family
```

© 2026 IP Infusion Inc. Proprietary

BGP Flowspec Commands

This section contains the BGP Flowspec commands.

address-family 1441
address-family ipv4 flowspec vrf 1442
class-map type traffic match-all 1443
drop 1444
flowspec 1445
hardware filter ipv4-bgp-flowspec 1446
ipv4 flowspec-disable 1447
local-install 1451
match destination-address 1452
match destination-port 1453
match dscp 1454
match fragment-type 1455
match icmp-code 1456
match icmp-type 1457
match packet-length 1458
match port 1459
match protocol 1460
match source-address 1461
match source-port 1462
match tcp-flag 1463
neighbor flowspec-vaildation-disable 1464
police rate 1465
policy-map 1466
policy-map class type traffic 1467
redirect 1468
service-policy type pbr 1469
set dscp 1470
set ip next-hop 1471
set ipv6 next-hop 1472
show hsl bgp-flowspec-rules fib-id 1473
show ip bgp flowspec 1475
vrf 1481

© 2026 IP Infusion Inc. Proprietary

address-family

Specify address-family type to attach a policy to a specific address-family. This command enters address family configuration mode.

Note: IPv6 address-family is not supported.

Command Syntax

```
address-family ipv4
no address-family
```

Parameters

None

Default

None

Command Mode

Flowspec mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#flowspec
(config-flowspec)#address-family ipv4
```

© 2026 IP Infusion Inc. Proprietary

address-family ipv4 flowspec vrf

Use this command to enter address-family IPv4 Flowspec VRF mode to configure IPv4 BGP Flowspec specific parameters in a VRF.

Command Syntax

```
address-family ipv4 flowspec vrf VRFNAME
no address-famlily ipv4 flowspec vrf VRFNAME
```

Parameters

VRFNAME
VRF name

Default

None

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config-router)#address-family ipv4 flowspec vrf VRF1
```

© 2026 IP Infusion Inc. Proprietary

class-map type traffic match-all

Use this command to define a traffic class and associate rules by using match sub-commands, after the class-map type traffic command is typed in. The created c class-map can be removed with No form of this command. Only match-all criteria (match all match items configured in a class-map) is supported in BGP Flowspec.

Command Syntax

```
class-map type traffic match-all NAME
no class-map type traffic NAME
```

Parameters

NAME
A class-map-name to define.

Default

None

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#class-map type traffic match-all test1
```

drop

When an attached class-map is matched, the matched packet drops.

Command Syntax

```
drop
no drop
```

Parameters

None

Default

None

Command Mode

policy-map class mode.

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config-pmap-pbr-c)#drop
```

© 2026 IP Infusion Inc. Proprietary

flowspec

Use this command to enable flowspec in the global configuration mode.

Command Syntax

```
flowspec
no flowspec
```

Parameters

None

Default

None

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#flowspec
(config-flowspec)#
```

© 2026 IP Infusion Inc. Proprietary

hardware filter ipv4-bgp-flowspec

Use this command to enable ipv4 BGP flowspec feature in hardware. Unless it is enabled, the all flowspec routes are not set to hardware, while the all BGP operations are done not related to this command.

Command Syntax

```
hardware filter ipv4-bgp-flowspec enable
hardware filter ipv4-bgp-flowspec disable
```

Parameters

None

Default

None

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#hardware-profile filter ipv4-bgp-flowspec enable
```

© 2026 IP Infusion Inc. Proprietary

ipv4 flowspec-disable

Use this command to disable BGP FlowSpec on a per-port basis. This is achieved by installing a high-priority FlowSpec rule that matches a specific FIB ID and Layer 3 (L3) interface, but with no action defined. This rule takes precedence over regular FlowSpec rules, ensuring that packets received on the specified L3 interface do not match any other FlowSpec rules. As a result, the FlowSpec functionality is effectively disabled for those specific ports. This feature is useful for environments where certain ports need to be excluded from BGP FlowSpec rule enforcement while maintaining regular network operations.

Notes:

This feature can be applied to any L3 interfaces, specifically on the following types:

- Physical interfaces: Direct physical network interfaces.
- Link Aggregation Group (LAG) interfaces: Bundled multiple network connections to act as a single interface.
- Virtual Local Area Network (VLAN) interfaces: Logical interfaces representing VLANs.
- Sub-interfaces: Logical interfaces created on a physical or VLAN interface for managing traffic separately.

Use no form of this command to enable BGP FlowSpec on a per-port basis.

Command Syntax

```
ipv4 flowspec-disable
no ipv4 flowspec-disable
```

Parameters

None

Default

None

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS version 6.6.0.

Examples

In the following scenario, define traffic classes and policy maps to handle traffic flows. Enable FlowSpec and apply the policy maps to control traffic behavior. When disabling FlowSpec on VLAN and Port-Channel interfaces (`vlan1.100` and `po1`), the system installs disabling rules with higher priority for the interfaces `vlan1.100` and `po1`. This ensures that regular FlowSpec rules do not apply to `vlan1.100` and `po1`.

© 2026 IP Infusion Inc. Proprietary

Disabling FlowSpec on VLAN and Port-Channel Interfaces

```
OcNOS(config)#interface po1
OcNOS(config-if)#ip address 100.1.1.4/24
OcNOS(config-if)#ipv4 flowspec-disable
OcNOS(config-if)#commit
OcNOS(config-if)#exit

OcNOS(config)#interface vlan1.100
OcNOS(config-if)#ip address 200.1.1.4/24
OcNOS(config-if)#ipv4 flowspec-disable
OcNOS(config-if)#commit
OcNOS(config-if)#end
```

Running Configuration

```
OcNOS#show running-config
!
! Software version: EC_AS5912-54X-XP-6.5.3 09/19/2024 09:56:53
!
! Last configuration change at 05:43:39 UTC Tue Sep 24 2024 by ocnos
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
no feature netconf-ssh
no feature netconf-tls
service password-encryption
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile filter ipv4-bgp-flowspec enable
hardware-profile statistics ingress-acl enable
!
qos enable
!
hostname R4
no ip domain-lookup
ip domain-lookup vrf management
bridge 1 protocol rstp vlan-bridge
tfo Disable
errdisable cause stp-bpdu-guard
no feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
ntp enable vrf management
!
class-map type traffic match-all clas1
 match destination-address ipv4 100.1.1.0/24
!
class-map type traffic match-all clas2
 match destination-address ipv4 200.1.1.0/24
!
policy-map type pbr pol1
 class type traffic clas1
  drop
!
policy-map type pbr pol2
```

```
 class type traffic clas2
  drop
!
flowspec
 local-install interface-all
 address-family ipv4
  service-policy type pbr pol1
  service-policy type pbr pol2
 !
 exit
!
vlan database
 vlan 100-1000 bridge 1
!
ip vrf management
!
interface po1
 ip address 100.1.1.4/24
 ipv4 flowspec-disable
!
interface ce49
!
interface ce50
 channel-group 1 mode active
!
interface ce51
 channel-group 1 mode active
!
interface ce52
!
interface ce53
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan add 100
!
interface ce54
!
interface eth0
 ip vrf forwarding management
 ip address 10.14.110.108/24
!
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 vlan1.100
 ip address 200.1.1.4/24
 ipv4 flowspec-disable
!
 exit
!
router bgp 100
!
ip route vrf management 0.0.0.0/0 10.14.110.1 eth0
!
!
end
```

Validation

To verify the installation of FlowSpec rules, use the `show hsl bgp-flowspec-rules fib-id all detail` command.

1. When traffic arrives at `vlan1.100` or `po1`, the system first matches the high-priority disabling rules (rules 1 and 2). These rules prevent further FlowSpec rules from being applied to traffic on the VLAN and Port-Channel interfaces.
2. For other interfaces, the system applies the regular FlowSpec rules. If the traffic matches the criteria defined in the class-map type `clas1` or `clas2` (rules 3 and 4), it will be dropped.

The prioritization of the disabling rules for `vlan1.100` and `po1` ensures these interfaces are effectively excluded from FlowSpec processing.

```
OcNOS#show hsl bgp-flowspec-rules fib-id all detail

<<< HSL BGP Flowspec information >>>
 rule processing type: all rules install
 Table Full: FALSE

<<< fib ID: 0 >>>
  Total number of rules installed: 4
  Rules installed to hw: 4
  Ruules not installed to hw: 0
  Table Full: FALSE

  <<< rule number: 1 >>>
    priority: 1
    disabled_interface: vlan1.100 (ifindex:10205)

    <<< Traffic Filters >>>

    <<< Actions >>>

  <<< rule number: 2 >>>
    priority: 2
    disabled_interface: po1 (ifindex:100001)

    <<< Traffic Filters >>>

    <<< Actions >>>

  <<< rule number: 3 >>>
    priority: 3

    <<< Traffic Filters >>>
    destination ip address: 100.1.1.0/24
    name: clas1

    <<< Actions >>>
    traffic_rate: 0 pps

  <<< rule number: 4 >>>
    priority: 4

    <<< Traffic Filters >>>
    destination ip address: 200.1.1.0/24
    name: clas2

    <<< Actions >>>
    traffic_rate: 0 pps
```

© 2026 IP Infusion Inc. Proprietary

local-install

Specify whether or not, the attached service policy is installed to a local node. By default, not installed. This command can be taken effect for all address-families by specifying it before entering address-family configuration mode, or for a specific address-family when the mode is address-family configuration.

Command Syntax

```
local-install interface-all
no local-install
```

Parameters

None

Default

None

Command Mode

Flowspec, Flowspec address family and Flowspec-VRF mode.

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#flowspec
(config-flowspec)#local-install interface-all
(config-flowspec)#address-family ipv4
(config-flowspec-af)#local-install interface-all
(config-flowspec-af)#exit
(config-flowspec)#vrf VRF1
(config-flowspec-vrf)#local-install interface-all
```

© 2026 IP Infusion Inc. Proprietary

match destination-address

IP destination address to match. It can be specified with prefix and netmask bits.

Command Syntax

```
match destination-address ipv4 A.B.C.D/M
no match destination-address
```

Parameters

A.B.C.D/M
A.B.C.D is prefix and M is netmask bits.

Default

None

Command Mode

Class-map mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#class-map type traffic match-all test1
(config-cmap-tr)#match destination-address ipv4 1.1.1.1/32
```

© 2026 IP Infusion Inc. Proprietary

match destination-port

Use this command to match TCP or UDP destination port number(s).

Command Syntax

```
match destination-port WORD
no match destination-port
```

Parameters

WORD

A port number in 1-65535. A range in a form of range: MIN_VALUE-MAX_VALUE. e.g. 10-100

Default

None

Command Mode

Class-map mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#class-map type traffic match-all test1
(config-cmap-tr)#match destination-port 10-100
```

© 2026 IP Infusion Inc. Proprietary

match dscp

Use this command to match DSCP value.

Command Syntax

```
match dscp WORD
no match dscp
```

Parameters

WORD

A class-map-name to define DSCP code value. A range in a form of range: MIN_VALU-MAX_VALUE. e.g. 10-3e

Default

None

Command Mode

Class-map mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#class-map type traffic match-all test1
(config-cmap-tr)#match dscp 21
```

match fragment-type

Use this command to match fragment type.

Command Syntax

```
match fragment-type (df|isf|ff|lf|lf-ff-isf)
no match fragment-type
```

Parameters

df
Don't fragment

isf
Is s fragment other than the first

ff
First fragment

lf
Last fragment

lf-ff-isf
When one of lf, ff, and isf is true to match, it matches.

Default

None

Command Mode

Class-map mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#class-map type traffic match-all test1
(config-cmap-tr)#match fragment-type lf
```

© 2026 IP Infusion Inc. Proprietary

match icmp-code

Use this command to match the ICMP code.

Command Syntax

```
match ipv4 icmp-code WORD
no match ipv4 icmp-code
```

Parameters

WORD

ICMP code number in 0-255. A range in a form of range: MIN_VALU-MAX_VALUE. e.g. 10-100

Default

None

Command Mode

Class-map mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#class-map type traffic match-all test1
(config-cmap-tr)#match ipv4 icmp-type 10
(config-cmap-tr)#match ipv4 icmp-code 20
```

© 2026 IP Infusion Inc. Proprietary

match icmp-type

Use this command to match TCP/UDP destination port number(s).

Command Syntax

```
match ipv4 icmp-type WORD
no match ipv4 icmp-type
```

Parameters

WORD
ICMP code number in 0-255. A range in a form of range: MIN_VALU-MAX_VALUE. e.g. 10-100

Default

None

Command Mode

Class-map mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#class-map type traffic match-all test1
(config-cmap-tr)#match ipv4 icmp-type 10
```

match packet-length

Use this command to match packet length.

Command Syntax

```
match packet-length WORD
no match packet-length
```

Parameters

WORD

ICMP code number in 0-255 (packet length value). A range in a form of range: MIN_VALU-MAX_VALUE. e.g. 10-100

Default

None

Command Mode

Class-map mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#class-map type traffic match-all test1
(config-cmap-tr)#match packet-length 70-100
```

match port

Use this command to match TCP/UDP port number(s).

Command Syntax

```
match port WORD
no match port
```

Parameters

WORD

A port number in 1-65535. A range in a form of range: MIN_VALUE-MAX_VALUE. e.g. 10-100

Default

None

Command Mode

Class-map mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#class-map type traffic match-all test1
(config-cmap-tr)#match port 1000-2000
```

match protocol

Use this command to match Protocol number(s).

Command Syntax

```
match protocol WORD
no match protocol
```

Parameters

WORD

IP protocol number in 0-255. A range in a form of range: MIN_VALUE-MAX_VALUE. e.g. 10-100

Default

None

Command Mode

Class-map mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#class-map type traffic match-all test1
(config-cmap-tr)#match protocol 122
```

match source-address

IP source address to match. It can be specified with prefix and netmask bits.

Command Syntax

```
match source-address ipv4 A.B.C.D/M
no match source-address
```

Parameters

A.B.C.D/M
A.B.C.D is prefix and M is netmask bits.

Default

None

Command Mode

Class-map mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#class-map type traffic match-all test1
(config-cmap-tr)#match source-address ipv4 2.2.2.2/32
```

match source-port

Use the command to match TCP/UDP source port number(s).

Command Syntax

```
match source-port WORD
no match source-port
```

Parameters

WORD

A port number in 1-65535. A range in a form of range: MIN_VALU-MAX_VALUE. e.g. 10-100

Default

None

Command Mode

Class-map mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#class-map type traffic match-all test1
(config-cmap-tr)#match source-port 10-100
```

match tcp-flag

Use this command to match TCP flag.

Command Syntax

```
match tcp-flag WORD bit-mask WORD
no match tcp-flag bitmask
```

Parameters

tcp-flag WORD
TCP flag value in hexadecimal (0-fff)

bit-mask WORD
TCP flag bitmask in hexadecimal (0-fff)

Default

None

Command Mode

Class-map mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#class-map type traffic match-all test1
(config-cmap-tr)#match tcp-flag 2e bit-mask 3f
```

neighbor flowspec-vaildation-disable

The following parameters can be configured with neighbor command in address-family for BGP Flowspec. The flowspec-vaildation-disable is a parameter for BGP Flowspec. When it is specified, BGP Flowspec validation is disabled. By default, this parameter is enabled.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X) flowspec-validation-disable
```

Parameters

A.B.C.D
IPv4 address

X:X::X:X
IPv6 address

Default

Enabled

Command Mode

Router Address Family mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config-router)#address-family ipv4 flowspec
(config-router-af)#neighbor 1.1.1.1 flowspec-validation-disable
```

© 2026 IP Infusion Inc. Proprietary

police rate

Use this command to set a rate applied for matched packets.

Command Syntax

```
police rate <1-1000000000> (bps|pps)
no police rate
```

Parameters

police rate <1-1000000000>
Rate value specified in 1-1000000000.

pps
A unit for rate value, packets per second (pps).

bps
A unit for rate value, bit per second (bps).

Default

None

Command Mode

Policy-map class mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
OcNOS(config-pmap-pbr-c)#police rate
OcNOS(config-pmap-pbr-c)#police rate 100000000 bps
```

© 2026 IP Infusion Inc. Proprietary

policy-map

Create or modify a policy map with a pbr type that can be attached as a service policy, and proceed to a mode to enter class command after this command is typed in. Multiple class-map can be linked to a created policy-map. To delete a policy-map, use the no form of this command.

Command Syntax

```
policy-map type (pbr) NAME
no policy-map type pbr NAME
```

Parameters

pbr NAME
A name for a creating modifying policy-map.

Default

None

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#policy-map type pbr
(config)#policy-map type pbr test2
```

policy-map class type traffic

Use this command to attach a class to a policy-map created by policy-map command. This command enters policy-map mode.

Command Syntax

```
class type (traffic) NAME
no class type traffic NAME
```

Parameters

traffic NAME

A name for class-map already created with class-map command.

Default

None

Command Mode

Policy-map mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#policy-map type pbr test2
(config-pmap-pbr)#class type traffic test1
```

© 2026 IP Infusion Inc. Proprietary

redirect

Use this command to set a VRF redirected to for matched packets.

Command Syntax

```
redirect ipv4 extcommunity rt ASN:nn_or_IP-address:nn
no redirect ipv4 extcommunity rt
```

Parameters

ASN:nn_or_IP-address:nn

Specify a route target, Two formats can be used:

- AS number: index, e.g. 100:10
- IPv4 address: index, e.g. 2.2.2.2:10.

Default

None

Command Mode

Policy-map class mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config-pmap-pbr-c)#redirect ipv4 extcommunity rt 100:10
```

© 2026 IP Infusion Inc. Proprietary

service-policy type pbr

Use this command to attach or detach a service policy to BGP component. The attached policy is locally installed optionally by typing local-install command, as well as a BGP Flowspec route is injected to a BGP component.

Command Syntax

```
service-policy type (pbr) NAME
no service-policy type pbr NAME
```

Parameters

pbr NAME
Name of the policy map.

Default

None

Command Mode

Flowspec address family mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config)#flowspec
(config-flowspec)#local-install interface-all
(config-flowspec)#address-family ipv4
(config-flowspec-af)#service-policy type pbr test1
```

© 2026 IP Infusion Inc. Proprietary

set dscp

Use this command to set a DSCP value to a matched packet to mark with it.

Command Syntax

```
set dscp DSCP_VALUE
no set dscp
```

Parameters

DSCP_VALUE

DSCP value in hexadecimal (0-3f) -
(00|01|02|03|04|05|06|07|08|09|0a|0b|0c|0d|0e|0f|10|11|12|13|14|15|16|17|18|19|1a|1b|1c|1d|1e|1f|
20|21|22|23|24|25|26|27|28|29|2a|2b|2c|2d|2e|2f|30|31|32|33|34|35|36|37|38|39|3a|3b|3c|3d|3e|3f)

Default

None

Command Mode

Policy-map class mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
(config-pmap-pbr-c)#set dscp 3f
```

set ip next-hop

Use this command to configure ipv4 nexthop as tracked object to verify its availability.

Note: Depends on “IP-SLA” configurations.

Command Syntax

```
set ip next-hop A.B.C.D verify-availability track <1-500>
no set ip next-hop A.B.C.D
```

Parameters

A.B.C.D
Specifies the IP address of the next-hop router.

verify-availability
Enables availability tracking for the specified next-hop IP.

track <1-500>
Associates the route with a tracking object ID (in the range 1 to 500), which monitors the reachability or status of the next-hop. If the tracked object goes down, the next-hop is considered unavailable and the route is withdrawn or deprioritized.

Command Mode

policy-map class mode.

Applicability

This command was introduced in OcNOS version 6.6.1.

Examples

```
(config-pmap-pbr-c)#set ip next-hop 10.10.0.67 verify-availability track 7
```

set ipv6 next-hop

Use this command to configure IPv6 nexthop as tracked object to verify its availability.

Note: Depends on “IP-SLA” configurations.

Command Syntax

```
set ipv6 next-hop X:X::X:X verify-availability track <1-500>
no set ipv6 next-hop X:X::X:X
```

Parameters

X:X::X:X
Specifies the IPv6 address of the next-hop router.

verify-availability
Enables availability tracking for the specified next-hop IP.

track <1-500>
Associates the route with a tracking object ID (in the range 1 to 500), which monitors the reachability or status of the next-hop. If the tracked object goes down, the next-hop is considered unavailable and the route is withdrawn or deprioritized.

Command Mode

Policy-map class mode.

Applicability

This command was introduced in OcNOS version 6.6.1.

Examples

```
(config-pmap-pbr-c)#set ipv6 next-hop 1311::2 verify-availability track 8
```

show hsl bgp-flowspec-rules fib-id

Use these commands to display the BGP FlowSpec rules for FIB IDs.

Command Syntax

```
show hsl bgp-flowspec-rules fib-id (<0-255>|all) (detail|)
```

Parameters

fib-id <0-255>
Specifies the FIB ID range.

fib-id all
Displays rule for all the FIB IDs.

fib-id detail
(Optional) Displays the detailed information of the FlowSpec rules.

Default

None

Command Mode

Execution mode

Applicability

Introduced in OcNOS version 6.3.0.

Examples

```
OcNOS#show hsl bgp-flowspec-rules fib-id all detail

<<< HSL BGP Flowspec information >>>
rule processing type: all rules install
Table Full: FALSE

<<< fib ID: 0 >>>
  Total number of rules installed: 6
  Rules installed to hw: 6
  Ruules not installed to hw: 0
  Table Full: FALSE

  <<< rule number: 1 >>>
    priority: 1
    disabled_interface: vlan1.100 (ifindex:10154)

    <<< Traffic Filters >>>

    <<< Actions >>>

  <<< rule number: 2 >>>
    priority: 2
    disabled_interface: vlan1.200 (ifindex:25200)

    <<< Traffic Filters >>>
```

© 2026 IP Infusion Inc. Proprietary

```
    <<< Actions >>>

  <<< rule number: 3 >>>
    priority: 3

    <<< Traffic Filters >>>
    destination ip address: 100.1.1.0/24
    name: clas1
    <<< Actions >>>
    traffic_rate: 0 pls

  <<< rule number: 4 >>>
    priority: 4

    <<< Traffic Filters >>>
    destination ip address: 200.1.1.0/24
    name: clas2

    <<< Actions >>>
    traffic_rate: 0 pps

  <<< rule number: 5 >>>
    priority: 5
    <<< Traffic Filters >>>
    destination ip address: 200.2.1.0/24
    name: clas3

    <<< Actions >>>
    traffic_rate: 0 pps
  <<< rule number: 6 >>>
    priority: 6
    <<< Traffic Filters >>>
    destination ip address: 200.3.1.0/24
    name: clas4
    <<< Actions >>>
    traffic_rate: 0
```

show ip bgp flowspec

Use these commands to display the BGP Flowspec summary details.

Command Syntax

```
show ip bgp flowspec
show ip bgp flowspec summary (vrf (VRFNAME|all|default))
show ip bgp (ipv4|vpnv4) flowspec (detail|)
show ip bgp (ipv4|vpnv4) flowspec (vrf (VRFNAME|all|default)) (detail|)
show ip bgp vpnv4 flowspec rd WORD (detail|)
```

Parameters

vrf VRFNAME
Display the Flowspec specific VRF details

all
Display the Flowspec all VRF details

default
Display the Flowspec default VRF details

detail
Display all Flowspec details

Default

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
#show ip bgp flowspec
BGP router identifier 19.19.19.19, local AS number 100
BGP table version is 7
1 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   40         35       7      0      0  00:04:51               1

Total number of neighbors 1

Total number of Established sessions 1
BGP router identifier 4.4.4.4, local AS number 100
BGP VRF VRF1 Route Distinguisher: 100:10
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
200.2.1.2               4 65530    0          0       0      0      0     never     Connect

Total number of neighbors 1

Total number of Established sessions 0
```

```
#show ip bgp flowspec summary vrf VRF1
BGP router identifier 4.4.4.4, local AS number 100
BGP VRF VRF1 Route Distinguisher: 100:10
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
200.2.1.2               4 65530    0          0       0      0      0     never        Idle

Total number of neighbors 1

Total number of Established sessions 0
```

#show ip bgp flowspec summary vrf all

```
#show ip bgp flowspec summary vrf all
BGP router identifier 4.4.4.4, local AS number 100
BGP VRF VRF1 Route Distinguisher: 100:10
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
200.2.1.2               4 65530    0          0       0      0      0     never     Connect

Total number of neighbors 1

Total number of Established sessions 0
BGP router identifier 19.19.19.19, local AS number 100
BGP table version is 7
1 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   41         37       7      0      0  00:05:31              1

Total number of neighbors 1

Total number of Established sessions 1
```

```
#show ip bgp flowspec summary vrf default
BGP router identifier 19.19.19.19, local AS number 100
BGP table version is 7
1 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   42         37       7      0      0  00:05:37              1

Total number of neighbors 1

Total number of Established sessions 1
```

```
#show ip bgp ipv4 flowspec
```

```
detail   summary  vrf
#show ip bgp ipv4 flowspec vrf all
BGP table version is 1, local router ID is 4.4.4.4
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

BGP Route Table for VRF VRF1
*>i  flowspec/368     1.1.1.1              0        100       0    i

Total number of prefixes 1
*>i  flowspec/368     1.1.1.1              0        100       0    i

Total number of prefixes 1
```

```
#show ip bgp ipv4 flowspec vrf all detail
BGP table version is 1, local router ID is 4.4.4.4
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

BGP Route Table for VRF VRF1
*>i  flowspec/368     1.1.1.1              0        100       0    i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 2.2.2.2/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:122
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1000
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:2000
    ICMP type:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:20
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x2e
      e:1 a:1 len:0 not:1 m:0 data:0x11
    Packet length:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:70
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:100
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x1c
    fragment:
      e:1 a:0 len:0 not:0 m:1 data:0x8

  FlowSpec Actions:
    traffic_rate: 100000000 bps
    rt_redirect: as_type: as 100:10 vrf_id:2 fib_id:2


Total number of prefixes 1
*>i  flowspec/368     1.1.1.1              0        100       0    i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 2.2.2.2/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:122
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1000
```

```
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:2000
    ICMP type:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:20
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x2e
      e:1 a:1 len:0 not:1 m:0 data:0x11
    Packet length:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:70
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:100
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x1c
    fragment:
      e:1 a:0 len:0 not:0 m:1 data:0x8

  FlowSpec Actions:
    traffic_rate: 0 pps


Total number of prefixes 1
```

```
#show ip bgp vpnv4 flowspec vrf VRF1
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
Route Distinguisher: 100:10 (Default for VRF VRF1)
*>i  flowspec/368     1.1.1.1               0        100        0    i
 Announced routes count = 0
 Accepted routes count = 1
#show ip bgp vpnv4 flowspec vrf VRF1 detail
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
Route Distinguisher: 100:10 (Default for VRF VRF1)
*>i  flowspec/368     1.1.1.1               0        100        0    i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 2.2.2.2/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:122
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1000
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:2000
    ICMP type:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:20
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x2e
      e:1 a:1 len:0 not:1 m:0 data:0x11
    Packet length:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:70
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:100
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x1c
    fragment:
      e:1 a:0 len:0 not:0 m:1 data:0x8
```

```
  FlowSpec Actions:
    traffic_rate: 100000000 bps
    rt_redirect: as_type: as 100:10 vrf_id:2 fib_id:2

 Announced routes count = 0
 Accepted routes count = 1
```

```
#show ip bgp vpnv4 flowspec rd 100:10
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
Route Distinguisher: 100:10 (Default for VRF VRF1)
*>i  flowspec/368     1.1.1.1               0        100        0    i
 Announced routes count = 0
 Accepted routes count = 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
Route Distinguisher: 100:10
*>i  flowspec/368     1.1.1.1               0        100        0    i
 Announced routes count = 0
 Accepted routes count = 1
```

```
#show ip bgp vpnv4 flowspec rd 100:10 detail
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
Route Distinguisher: 100:10 (Default for VRF VRF1)
*>i  flowspec/368     1.1.1.1               0        100        0    i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 2.2.2.2/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:122
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1000
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:2000
    ICMP type:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:20
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x2e
      e:1 a:1 len:0 not:1 m:0 data:0x11
    Packet length:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:70
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:100
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x1c
    fragment:
      e:1 a:0 len:0 not:0 m:1 data:0x8

  FlowSpec Actions:
    traffic_rate: 100000000 bps
```

```
    rt_redirect: as_type: as 100:10 vrf_id:2 fib_id:2

 Announced routes count = 0
 Accepted routes count = 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
Route Distinguisher: 100:10
*>i  flowspec/368     1.1.1.1               0        100        0    i

  FlowSpec:
    destination ip address: 1.1.1.1/32
    source ip address: 2.2.2.2/32
    IP protocol:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:122
    TCP/UDP port:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:1000
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:2000
    ICMP type:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:10
    ICMP code:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:20
    TCP flags:
      e:0 a:0 len:0 not:0 m:1 data:0x2e
      e:1 a:1 len:0 not:1 m:0 data:0x11
    Packet length:
      e:0 a:0 len:1 lt:0 gt:1 eq:1 data:70
      e:1 a:1 len:1 lt:1 gt:0 eq:1 data:100
    DSCP:
      e:1 a:0 len:0 lt:0 gt:0 eq:1 data:0x1c
    fragment:
      e:1 a:0 len:0 not:0 m:1 data:0x8

  FlowSpec Actions:
    traffic_rate: 100000000 bps
    rt_redirect: as_type: as 100:10 vrf_id:2 fib_id:2

 Announced routes count = 0
 Accepted routes count = 1
```

vrf

Use this command to specify a VRF to configure service policies. This command is executed in the default VRF mode, and enter VRF mode. address-family, local-install and service-policy commands can be executed in the same way (stated in `address-family, locall-install,` and `service-policy type pbr`) executed in the default VRF mode.

Command Syntax

```
vrf VRF_NAME
no vrf VRF_NAME
```

Parameters

VRF_NAME
Specify a VRF name already created

Default

None

Command Mode

Flowspec mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
OcNOS(config)#flowspec
OcNOS(config-flowspec)#vrf VRF1
```

BGP Virtual Private Network Commands

This section describes the BGP Virtual Private Network (VPN) configuration commands.

Notes: For the following commands refer to the BGP Commands (page 1244) and BGP4+ Commands (page 1423) sections

- address-family ipv4 (page 1248)
- address-family ipv6 (page 1424)
- exit-address-family (page 1319)
- neighbor activate (page 1338)
- neighbor allowas-in (page 1340)
- neighbor as-origination-interval (page 1342)
- neighbor description (page 1356)
- neighbor remote-as (page 1387)
- neighbor send-community (page 1396)
- neighbor shutdown (page 1399)
- redistribute (page 1415)

bgp external-route-leak 1483
bgp inbound-route-filter 1484
clear ip bgp <1-4294967295> vpnv4 1485
clear ip bgp A.B.C.D vpnv4 1486
clear bgp <1-4294967295> l2vpn vpls 1487
clear bgp A.B.C.D l2vpn vpls 1488
clear bgp * l2vpn vpls 1489
clear ip bgp * vpnv4 1490
debug bgp mpls 1491
export map 1492
import map 1495
ip vrf 1496
neighbor allow-ebgp-vpn 1497
neighbor as-override 1498
neighbor send-community 1499
neighbor soo 1500
rd (route distinguisher) 1501
route-target 1502

© 2026 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
```

Parameter

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. Inbound-route-filter support for L2VPN EVPN introduced from OcNOS version 6.3.0 .

Examples

```
#configure terminal
(config)#router bgp 100
(config-router)#bgp inbound-route-filter
```

© 2026 IP Infusion Inc. Proprietary

clear ip bgp <1-4294967295> vpnv4

Use this command to reset a BGP connection for all VPN peers in a specified Autonomous System.

Note: The `soft in` or `in` and `soft out` or `out` in the BGP commands performs the same functionality. User can use any of the commands to soft reset.

Command Syntax

```
clear ip bgp <1-4294967295> vpnv4 unicast in
clear ip bgp <1-4294967295> vpnv4 unicast out
clear ip bgp <1-4294967295> vpnv4 unicast soft
clear ip bgp <1-4294967295> vpnv4 unicast soft in
clear ip bgp <1-4294967295> vpnv4 unicast soft out
```

Parameters

<1-4294967295>
Clear peers with this AS number

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 execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip bgp 500 vpnv4 unicast soft out
```

clear ip bgp A.B.C.D vpnv4

Use this command to reset an VPNv4 BGP connection for a specific IPv4 address.

Note: The `soft in` or `in` and `soft out` or `out` in the BGP commands performs the same functionality. User can use any of the commands to soft reset.

Command Syntax

```
clear ip bgp A.B.C.D vpnv4 unicast in
clear ip bgp A.B.C.D vpnv4 unicast out
clear ip bgp A.B.C.D vpnv4 unicast soft
clear ip bgp A.B.C.D vpnv4 unicast soft in
clear ip bgp A.B.C.D vpnv4 unicast soft out
```

Parameters

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 execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip bgp 10.10.0.12 vpnv4 unicast soft
#clear ip bgp 10.10.0.10 vpnv4 unicast out
```

© 2026 IP Infusion Inc. Proprietary

clear bgp <1-4294967295> l2vpn vpls

Use this command to reset the session for the neighbors with a specific ASN number for L2VPN VPLS.

Command Syntax

```
Clear bgp <1-4294967295> l2vpn vpls
```

Parameters

<1-4294967295>
Autonomous System number of the BGP neighbor.

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear bgp 100 l2vpn vpls
```

© 2026 IP Infusion Inc. Proprietary

clear bgp A.B.C.D l2vpn vpls

Use this command to reset the session for neighbor with address A.B.C.D.

Command Syntax

```
clear bgp A.B.C.D l2vpn vpls
```

Parameters

A.B.C.D
BGP neighbor address.

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear bgp 192.168.0.3 l2vpn vpls
```

clear bgp * l2vpn vpls

Use this command to reset the session with all neighbors for VPLS address family

Command Syntax

```
clear bgp * l2vpn vpls
```

Parameters

None

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear bgp * l2vpn vpls
```

© 2026 IP Infusion Inc. Proprietary

clear ip bgp * vpnv4

Use this command to reset a VPNv4 BGP connection for all peers. This command clears the BGP connection and dynamically resets the outbound routing table. This frees up additional memory required for storing updates to generate new updates.

Note: The `soft in` or `in` and `soft out` or `out` in the BGP commands performs the same functionality. User can use any of the commands to soft reset.

Command Syntax

```
clear ip bgp * vpnv4 unicast in
clear ip bgp * vpnv4 unicast out
clear ip bgp * vpnv4 unicast soft
clear ip bgp * vpnv4 unicast soft in
clear ip bgp * vpnv4 unicast soft out
```

Parameters

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 execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip bgp *
#clear ip bgp * vpnv4 unicast out
```

© 2026 IP Infusion Inc. Proprietary

debug bgp mpls

Use this command to enable the display of MPLS related information.

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

Note: This command is available only when `vrf` option is enabled.

Command Syntax

```
debug bgp mpls
no debug bgp mpls
```

Parameters

None

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
debug bgp mpls
```

export map

Use this command in VRF configuration mode to apply a route-map that controls which routes are exported from the VRF into the global routing table or into other VRFs. The route-map can filter routes using match clauses and modify BGP attributes using set clauses, enabling granular control over inter-VRF route export policies.

Notes:

- An export-map with a `permit route-map` that contains `no match` or `set clauses` functions as a **permit-all**.
- An export-map with a `deny route-map` that contains `no match` or `set clauses` functions as a **deny-all**.

Use the `no` parameter of this command to remove the applied route-map.

Command Syntax

```
export map WORD
no export map
```

Parameters

WORD

Specifies the name of the route-map to apply as the VRF export-map.

- `match` clauses filter the routes to export (for example, using prefix lists, community values, or AS path).
- `set` clauses adjust BGP attributes for exported routes (for example, route targets, local preference, AS path, MED).

Default

If no export-map is configured, all eligible VRF routes are exported unmodified.

Command Mode

VRF mode

Applicability

Introduced in OcNOS version 4.1. OcNOS version 7.0.0 adds support for VRF-level BGP export-maps with route-map match and set clauses, enabling granular control of exported routes and attributes.

Examples

This example shows how to configure an export-map with a route-map that filters specific prefixes and modifies their BGP attributes before exporting them from the VRF.

- The VRF `vrf1` is created, with route-distinguisher 1:1 and route-target 100:1.
- The `export map export-r1` is applied to control which routes `vrf1` exports.
- The prefix-list `p1_ipv4` permits only `121.1.1.1/32`.

© 2026 IP Infusion Inc. Proprietary

- The `route-map export-r1` matches the permitted prefix and modifies its attributes by:
  - Setting local preference to `1000`
  - Setting BGP community to `650:650`
  - Adding extended communities (route-targets) `110:1 and 220:1`

As a result, only the permitted prefix (`121.1.1.1/32`) is exported from `vrf1` with the modified BGP attributes.

```
OcNOS(config)#ip vrf vrf1
OcNOS(config-vrf)#export map export-r1
OcNOS(config-vrf)#
```

Prefix-list Definition

```
OcNOS#show running-config prefix-list
!
ip prefix-list p1_ipv4
 seq 5 permit 121.1.1.1/32
!
```

Route-map Definition

```
OcNOS#show running-config route-map
!
route-map export-r1 permit 10
 match ip address prefix-list p1_ipv4
 set local-preference 1000
 set community 650:650
 set extcommunity rt 110:1 220:1 additive
!
```

VRF Configuration

```
OcNOS#show running-config vrf
!
ip vrf management
!
ip vrf vrf1
 rd 1:1
 route-target both 100:1
 export map export-r1
!
ip vrf vrf2
 rd 2:2
 route-target import 100:1
 route-target both 200:1
!
```

BGP Configuration

```
OcNOS#show running-config bgp
!
router bgp 200
 bgp auto-policy-soft-reset enable
 neighbor 5.5.5.5 remote-as 200
 neighbor 5.5.5.5 update-source lo
 !
 address-family vpnv4 unicast
 neighbor 5.5.5.5 activate
 exit-address-family
 !
 address-family ipv4 vrf vrf1
 neighbor 1.1.1.1 remote-as 100
 neighbor 1.1.1.1 activate
 neighbor 2.2.2.2 remote-as 300
 neighbor 2.2.2.2 activate
 exit-address-family
```

© 2026 IP Infusion Inc. Proprietary

```
 !
 address-family ipv4 vrf vrf2
 exit-address-family
 !
 exit
!
```

import map

Use this command to assign a route-map to the VRF. This route-map is applied to routing information imported from another PE or VRF.

Also, apply this command when finer control is required over the routes imported into a VRF than what is provided by the standard import and export extended communities.

The route-map can be used to filter specific routes that are eligible for import into the VRF. Routes that match a deny clause in the route-map are excluded, while those that match a permit clause are imported.

Note: Only match rules (permit or deny) in the route-map are applied for imported routes. Any set rules defined in the route-map are ignored.

Use the `no` parameter of this command to remove the assigned route-map.

Command Syntax

```
import map WORD
no import map
```

Parameters

WORD

Specifies the name of the route-map to be applied for import control.

Default

None

Command Mode

VRF mode

Applicability

Introduced in OcNOS version 4.1.

Examples

Applies the route-map `set-pref` to control which routes are imported into the VRF `myVRF`.

```
(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 (page 1411) 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-vpnv4 mode and 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)#address-family vpnv4 unicast
(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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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 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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1411) 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
```

© 2026 IP Infusion Inc. Proprietary

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.

Route-target values on a MAC-VRF can be manually configured or auto derived from BGP.

Command Syntax

```
route-target (import|export|both) (ASN:nn_or_IP-address:nn|evpn-auto-rt)
no route-target (import|export|both) (ASN:nn_or_IP-address:nn|evpn-auto-rt)
```

Parameters

import
Import routing information

export
Export routing information

Note: Users can configure up to 128 export route-targets under "IP VRF mode”.

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).

evpn-auto-rt
route-target auto-derived from BGP

Default

No default value is specified

Command Mode

VRF mode

Applicability

This command was introduced before OcNOS version 1.3 and `evpn-auto-rt` option for mac-vrf is introduced from in OcNOS version 6.0.0.

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

(config)#mac vrf l2vrf1
(config-vrf)#route-target both evpn-auto-rt
```

BGP Show Commands

This section describes the BGP show commands.

show bgp 1506
show bgp A.B.C.D 1507
show bgp A.B.C.D/M 1509
show bgp client 1511
show bgp community 1512
show bgp community-list 1514
show bgp dampening dampened-paths 1515
show bgp dampening flap-statistics 1517
show bgp dampening parameters 1519
show bgp filter-list 1521
show bgp inconsistent-as 1523
show bgp ipv6 1524
show bgp ipv6 peer-group 1527
show bgp l2vpn vpls 1528
show bgp label-pool 1531
show bgp neighbors 1532
show bgp neighbors advertised-routes 1536
show bgp neighbors received prefix-filter 1537
show bgp neighbors received-routes 1538
show bgp neighbors routes 1540
show bgp nexthop-tracking 1542
show bgp nexthop-tree-details 1544
show bgp paths 1545
show bgp prefix-list 1546
show bgp quote-regexp 1548
show bgp regexp 1549
show bgp route-map 1550
show bgp statistics 1552
show bgp summary 1554
show bgp X:X::X:X 1557
show bgp X:X::X:X/M longer prefixes 1558
show debugging bgp 1559
show ip bgp 1560
show ip bgp cidr-only 1563
show ip bgp community-info 1564
show ip bgp peer-group 1565

© 2026 IP Infusion Inc. Proprietary

show ip bgp peer-group vrf all 1566
show ip bgp rtfilter all 1567
show ip bgp rtfilter neighbors 1570
show ip bgp scan 1572
show ip bgp vpnv4 1573
show ip bgp vpnv6 all neighbors 1579
show ip bgp vpnv6 rd neighbors 1582
show ip extcommunity-list 1584
show ip protocols 1585
show ip vrf 1586
show running-config as-path access-list 1587
show running-config bgp 1588
show running-config community-list 1589
show running-config vrf 1590
show running-config ip vrf 1591
show running-config mac vrf 1592
show running-config vrf word 1593

© 2026 IP Infusion Inc. Proprietary

show bgp

Use this command to display the status of BGP routes.

Command Syntax

```
show bgp ipv4
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

Execution modePrivilege 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

Execution modePrivilege mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#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 17 explains the output fields.

Table 17. 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

© 2026 IP Infusion Inc. Proprietary

Table 17. show ip bgp output details (continued)

Field	Description
	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 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

Execution modePrivilege 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
```

© 2026 IP Infusion Inc. Proprietary

```
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

Execution modePrivilege mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
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#
```

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege mode

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege 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]
```

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip bgp dampening dampened-paths
BGP table version is 32, 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

   Network           From                Reuse    Path
 d   11.11.11.0/24    90.90.90.1           00:27:20 200 i
```

Table 18 shows the status codes displayed at the start of a route entry.

Table 18. 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.

Table 19 shows the codes at the end of each route entry that indicate where the route originated.

Table 19. 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 20 explains the output fields.

Table 20. show bgp dampening dampened-paths output details

Field	Description
Network	Internet address of a network.
From	IP address of the advertising peer.
Reuse	The amount of time remaining until this route will be un-suppressed and can be used again.
Path	Autonomous system path to the destination network.

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege 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

Table 21 shows the status codes displayed at the start of a route entry.

Table 21. 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.

Table 22 shows the codes at the end of each route entry that indicate where the route originated.

Table 22. 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 23 explains the output fields.

Table 23. 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).
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.
Path	Autonomous system path to the destination network.

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege 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 24explains the output fields.

Table 24. 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).

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege mode

Applicability

This command was introduced before OcNOS version 1.3.

© 2026 IP Infusion Inc. Proprietary

Example

```
#show bgp filter-list bgp-local-onlyBGP 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-asshow 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

Execution modePrivilege 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      ?
```

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege 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 25 shows the status codes displayed at the start of a route entry.

Table 25. 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 26 shows the codes at the end of each route entry that indicate where the route originated.

Table 26. 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 27 explains the output fields.

Table 27. 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.
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.
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.

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege 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
#
```

© 2026 IP Infusion Inc. Proprietary

show bgp l2vpn vpls

Command Syntax

```
show bgp l2vpn vpls (rr|) (detail|)
```

Parameters

rr

Display the information of auto-discovered peers at Route reflector node.

detail

Display the detailed information of auto-discovered peers.

Command Mode

Execution modePrivilege mode

Applicability

This command was introduced in version OcNOS version 1.3.

Example

```
#show bgp l2vpn vpls
VPLS-ID   VE-ID    Discovered-Peers   Route-Target
10        3        1                   10:100


#show bgp l2vpn vpls  detail

VPLS ID: 10
VE-ID: 3
Discovered Peers: 1
Route-Target: 10:100
Local RD: 10:100
 Mesh Peers:
  Address:3.3.3.3, RD:10:100, VE-ID:4
  VC Details: VC-ID:34
  Remote (LB:52480,VBO:1,VBS:64)  Local (LB:52480,VBO:1,VBS:64)
  LB sent on known VEID:Yes
  In Label:52483, Out Label:52482
  PW Status:Established
```

Table 28 explains the output fields.

Table 28. show bgp l2vpn vpls output details

Field	Description
VPLS-ID	L2VPN address family database information for the Virtual Private LAN Service (VPLS).
VE-ID	L2VPN address family database information for the Virtual Expansion.
Discovered Peers	Peer discovery is used to find peers that are available for data using LLDP.

© 2026 IP Infusion Inc. Proprietary

Table 28. show bgp l2vpn vpls output details (continued)

Field	Description
Route-Target	An identifier prepended to IP addresses to assure the uniqueness of the address.
Local RD	The Local Route Descriptor – the first two numbers of the Route-Target.
Mesh Peers	Internal BGP peers – devices that do not re-advertise routes to other IBGP devices.
Address	Mesh session information for the peer specified with the ip-address argument.
RD	Mesh peer's Route-Descriptor.
VC Details	The virtual circuit session information with the ip-address for the Provider Edge (PE) routers.
Remote	LB (Label Base) – the first label value of a free set of labels that can be reserved by the PE router to be used for this VPLS domain. VBO (VE Block Offset) – the offset value to be used when multiple label blocks must be created by a PE router. VBS (VE Block Size) – the size of the label block.
Local	LB (Label Base) – the first label value of a free set of labels that can be reserved by the PE router to be used for this VPLS domain. VBO (VE Block Offset) – the offset value to be used when multiple label blocks must be created by a PE router. VBS (VE Block Size) – the size of the label block.
LB sent on known VEID	Whether the Label Base came on a known Virtual Expansion Identifier – yes or no.
In Label	The ingress (incoming interface) label for this segment.
Out Label	Label received from downstream neighbor for route.
PW Status	The status of the VPLS Pseudo-Wire. Values can be: Idle, Active, Open, or Established.

```
#show bgp l2vpn vpls rr
RD        RR-Clients     Non-Clients           Route-Target
10:100     2                  0                  10:100
10:100     2                  0                  10:100


#show bgp l2vpn vpls rr detail

Route-Target: 10:100
  Peer:1.1.1.1
    RR Client  : Yes
    VE-ID:3  LB:52480  VBO:1  VBS:64

Route-Target: 10:100
  Peer:3.3.3.3
    RR Client  : Yes
    VE-ID:4  LB:52480  VBO:1  VBS:64
```

Table 29 explains the output fields.

Table 29. show bgp l2vpn vpls rr output details

Field	Description
Route-Target	An identifier prepended to IP addresses to assure the uniqueness of the address.
Peer	Internal BGP peers – devices that do not re-advertise routes to other IBGP devices.
RR Client	Device is a client of the Route Reflector – yes or no.
VE-ID	L2VPN address family database information for the Virtual Expansion.
LB	LB (Label Base) – the first label value of a free set of labels that can be reserved by the PE router to be used for this VPLS domain.
VBO	VBO (VE Block Offset) – the offset value to be used when multiple label blocks must be created by a PE router.
VBS	VBS (VE Block Size) – the size of the label block.

© 2026 IP Infusion Inc. Proprietary

show bgp label-pool

Use this command to display the label management in BGP protocol.

Command Syntax

```
show bgp label-pool (block-id <0-1638>|)
```

Parameters

block-id
Block identifier of a block that is allocated to the protocol.

Command Mode

Execution modePrivilege mode

Applicability

This command was introduced in version OcNOS version 6.5.4.

Example

```
#show bgp label-pool
Module: BGP, Label range: 16-1048575, Current block: 78
+----------+----------+----------+---------------+---------------+---------------+
block_id label_min label_max usable_labels free_labels first_free_label
+----------+----------+----------+---------------+---------------+---------------+
78 49920 50559 640 96 50464
Total - blocks: 1, used-labels: 544, free-labels: 96
```

show bgp neighbors

Use this command to display information about BGP neighbor connections.

Command Syntax

```
show bgp neighbors
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)(advertise-routes|)
show ip bgp ipv4 (unicast|multicast) neighbors (A.B.C.D|X:X::X:X)
show ip bgp neighbors (A.B.C.D|X:X::X:X) (hold-time|keepalive-interval|connection-retrytime)
show ip bgp neighbors (A.B.C.D|X:X::X:X) (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

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

© 2026 IP Infusion Inc. Proprietary

update
Update messages

open
Open messages

keepalive
Keepalive messages

Command Mode

Execution modePrivilege mode

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
TCP MSS: (800), Advertise TCP MSS: (800), Send TCP MSS: (800),  Receive TCP MSS:
 (536)
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 30 explains the output fields.

Table 30. show bgp neighbor output details

Field	Description
BGP neighbor	BGP session information for the neighbor with the ip-address argument.

Table 30. show bgp neighbor output details (continued)

Field	Description
remote AS	Remote Autonomous system used to exchange exterior routing information between neighboring ASs.
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 35. 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

Table 30. show bgp neighbor output details (continued)

Field	Description
	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.
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.
TCP MSS	This value indicates the configured TCP MSS value using the `neighbor A.B.C.D tcp-mss <40-1440>` command. If the TCP MSS value is not configured explicitly, the system shows a value of 0.
Advertise TCP MSS	The host communicates this value during the TCP connection establishment phase, specifically during the TCP handshake. It is based on the network path's Maximum Transmission Unit (MTU) and any TCP options set on the host.
Send TCP MSS	This value represents the MSS for the connection, typically the lesser of the two MSS values exchanged.
Receive TCP MSS	This value reflects the MSS that the Kernel believes the client TCP stack uses for receiving data, indicating the largest TCP segment size the client can handle for incoming packets.
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) 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) advertised-routes
show ip bgp ipv4 (unicast|multicast) neighbors (A.B.C.D|X:X::X:X) advertised-routes
```

Parameters

A.B.C.D
IPv4 neighbor

X:X::X:X
IPv6 neighbor

ipv4
IPv4 addresses

multicast
Multicast prefixes

unicast
Unicast prefixes

vrf
VPN Routing/Forwarding instance name

Command Mode

Execution modePrivilege 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
```

© 2026 IP Infusion Inc. Proprietary

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) received prefix-filter
show ip bgp neighbors (A.B.C.D|X:X::X:X) 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

ipv4
IPv4 addresses

unicast
Unicast prefixes

multicast
Multicast prefixes

Command Mode

Execution modePrivilege 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
```

© 2026 IP Infusion Inc. Proprietary

show bgp neighbors received-routes

Use this command to view all the received routes, including all AFI/SAFI routes, from a BGP neighbour. Executing this CLI displays the received routes from accessing:

- BGP local RIB or
- BGP ADJ_IN RIB only if the 'soft-reconfiguration inbound' feature is supported.

Command Syntax

```
show ip bgp neighbors (A.B.C.D|X:X::X:X) received-routes (vrf (VRFNAME|all|default)|)
show ip bgp ipv4 (unicast|multicast) neighbors (A.B.C.D|X:X::X:X) received-routes
```

Parameters

A.B.C.D
IPv4 address

X:X::X:X
IPv6 address

ipv4
IPv4 addresses

unicast
Unicast prefixes

multicast
Multicast prefixes

VRFNAME
VPN Routing/Forwarding instance name

all
All VRFs

deault
Default VRF (Global routing)

Command Mode

Execution modePrivilege mode

Applicability

This command was introduced before OcNOS version 1.3. and updated in OcNOS version 6.3.1.

Example

```
OcNOS#show bgp neighbors 77.242.226.32 received-routes
% Inbound soft reconfiguration is not supported, displaying the received routes from accessing BGP
local RIB
For address family: IPv4 Unicast
BGP table version is 610, local router ID is 14.14.14.14
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,
l - labeled, S Stale
```

© 2026 IP Infusion Inc. Proprietary

```
Origin codes: i - IGP, e - EGP, ? - incomplete
Network Next Hop Metric LocPrf Weight Path
* i 20.20.20.0/24 77.242.226.32 0 100 0 ?
*>i 77.242.226.32/32 77.242.226.32 0 100 0 ?
*>i 200.10.1.0 77.242.226.32 0 100 0 ?
Total number of prefixes 3
For address family: VPNv4 Unicast
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
Route Distinguisher: 77.242.226.32:100
*>i 200.10.1.0 77.242.226.32 0 100 0 ?
Accepted routes count = 1
Route Distinguisher: 65535:22
*>i 10.50.100.160/30 10.10.40.133 0 100 0 ?
*>i 10.202.167.0/24 10.10.40.133 0 100 0 ?
Accepted routes count = 2

OcNOS#sh bgp neighbors 10.10.40.133 received-routes | begin VPNv4
...skipping
For address family: VPNv4 Unicast
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
Route Distinguisher: 65535:22
*>i 10.50.100.160/30 10.10.40.133 0 100 0 ?
*>i 10.202.167.0/24 10.10.40.133 0 100 0 ?
Accepted routes count = 2
To view the received-routes of particular address-family, use below filter.
“sh bgp neighbors 10.10.40.133 received-routes | begin LINE”
LINE can be of { IPv4 Unicast or
IPv4 Labeled-Unicast or
VPNv4 Unicast or
IPv6 Unicast or
6PE Labeled Unicast or
VPNv6 Unicast
}
```

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) routes
show ip bgp neighbors (A.B.C.D|X:X::X:X) routes
show ip bgp ipv4 (unicast|multicast) neighbors (A.B.C.D|X:X::X:X) routes
```

Parameters

A.B.C.D
IPv4 address

X:X::X:X
IPv6 address

ipv4
IPv4 addresses

unicast
Unicast prefixes

multicast
Multicast prefixes

Command Mode

Execution modePrivilege 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 31 explains the output fields.

Table 31. 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: 1. Route refresh 2. Address family IPv4 Unicast 3. 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.

© 2026 IP Infusion Inc. Proprietary

show bgp nexthop-tracking

Use this command to display BGP nexthop-tracking status.

Command Syntax

```
show bgp nexthop-tracking
```

Parameters

None

Command Mode

Execution modePrivilege 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 32 explains the output fields.

Table 32. 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.
BGP VRF	Name and ID number of this BGP VRF.

© 2026 IP Infusion Inc. Proprietary

Table 32. show bgp nexthop-tracking output details (continued)

Field	Description
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.

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege 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 33 explains the output fields.

Table 33. 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

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege 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 34 explains the output fields.

Table 34. 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.

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege 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
```

© 2026 IP Infusion Inc. Proprietary

```
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

Execution modePrivilege 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

Execution modePrivilege 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
```

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege 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
```

© 2026 IP Infusion Inc. Proprietary

```
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

Execution modePrivilege 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
MemPool - BGP MPLS REQ          : | Total (0/0) blk_size:32
#
```

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege 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 35 explains the fields for each neighbor entry.

Table 35. 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.
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:

Table 35. neighbor entry fields (continued)

Field	Description
	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.

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege 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

Execution modePrivilege 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

Execution modePrivilege 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
```

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege 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 36 explains the status codes in the header.

Table 36. 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 37 explains the codes are at the end of each routing entry that show where the route originated.

Table 37. 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 38 explains the fields shows for each route.

Table 38. 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.

© 2026 IP Infusion Inc. Proprietary

Table 38. route entry fields (continued)

Field	Description
	The status codes are explained in Table 36.
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 37

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).

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege 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
```

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip bgp community-info
Address Refcnt Community
[0x262a4c0] (1) 1:1
```

© 2026 IP Infusion Inc. Proprietary

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

WORD
Name of the peer group

Command Mode

Execution modePrivilege 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
```

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege 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
```

© 2026 IP Infusion Inc. Proprietary

show ip bgp rtfilter all

Use this command to display Route Target Constraint (RTC) Network Layer Reachability Information (NLRI) entries received from and advertised to BGP peers.

Command Syntax

```
show ip bgp rtfilter (all | summary)
```

Parameters

None

Command Mode

Execution mode and Privilege mode

Applicability

Introduced before OcNOS version 1.3.

Example

show ip bgp rtfilter all

```
#show ip bgp rtfilter all
RTFilter's Received
******************
peer-ip 22.22.22.22
100:1:400/96

RTFilter's Sent
******************
peer-ip 22.22.22.22
100:1:400/96

RTFilter's rt_add_pending
******************

RTFilter's rt_del_pending
******************

RTFilter's rtfilter_receive_pending
******************

 Example: The format 100:1:400/96 is interpreted as ORIGIN_AS:Route-Target/Prefix-length
 Breakdown of 100:1:400/96:
         The first part (100) is the Global Administrator ASN represents Origin AS.
         1:400 -> Represents the Route Target (RT) Extended Community.
         /96 -> The prefix length, which in the case of RTC is always 96 bits.
```

show ip bgp rtfilter all summary

```
#show ip bgp rtfilter all summary
BGP router identifier 192.168.36.11, local AS number 4294967294
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   Desc
192.168.36.1       4 4294967294        907      22670       1      0      0  00:26:24            43
192.168.36.2       4 4294967294        356       6046       0      0      0  01:23:25          Idle
192.168.36.3       4 4294967294        706      24519       1      0      0  00:25:49             6
192.168.36.4       4 4294967294       1015      21917       1      0      0  00:28:10          Idle
192.168.36.12      4 4294967294      12928       8239       1      0      0  00:27:32             2

Total number of neighbors 5

Total number of Established sessions 3
```

Here is the explanation of the show output fields.

Table 39. Description of output field for the 'show ip bgp rtfilter all' command.

Field	Description
RTFilter's Received	Lists the RTC NLRIs received from the specified BGP peer. These entries indicate the Route Targets (RTs) that the peer is interested in receiving.
RTFilter's Sent	Lists the RTC NLRIs advertised to the specified BGP peer. These entries indicate the Route Targets (RTs) for which this router is requesting VPN routes from the peer.
Peer-IP	Displays the IP address of the BGP peer associated with the RTC entries. **Example:** `peer-ip 22.22.22.22`
RTC NLRI Entry	The notation `100:1:400/96` in the output indicates a RTC NLRI entry, which is used in the Route Target Filtering mechanism of MPLS Layer 3 VPNs (L3VPNs). This limits the advertisement of VPN routes to only those needed by the peer. **Example:** Here is the breakdown of `100:1:400/96`: • 100 – Global Administrator ASN represents origin AS. • 1:400 – Route Target (RT) Extended Community. • /96 – Prefix length. For RTC NLRIs, this value is always 96 bits.
RTFilter's rt_add_pending	Shows RT filter routes that are pending to be added but not yet installed.
RTFilter's rt_del_pending	Shows RT filter routes that are pending to be deleted but not yet removed.
RTFilter's rtfilter_receive_pending	Shows RT filter routes that are pending to be processed after being received from peers.

Table 40. Description of output field for the 'show ip bgp rtfilter all summary' command.

Field	Description
BGP router identifier	The Router ID of the local BGP instance.
local AS number	The Autonomous System (AS) number configured on the local router.
BGP table version	Version number of the local BGP routing table.
BGP AS-PATH entries	Number of unique AS-PATH entries stored in the BGP table.
BGP community entries	Number of BGP community entries stored in the table.
Neighbor	IP address of the BGP peer.
V	BGP version in use with the neighbor.
AS	Autonomous System number of the neighbor.
MsgRcv	Total number of BGP messages received from the neighbor.
MsgSen	Total number of BGP messages sent to the neighbor.
TblVer	BGP table version that the neighbor is synchronized with.
InQ	Number of BGP messages currently in the input queue from the neighbor.
OutQ	Number of BGP messages currently in the output queue to the neighbor.
Up/Down	Elapsed time since the BGP session went up or down.
State/PfxRcd	If the session is established: displays the number of prefixes received from the neighbor. If the session is not established: shows the current state (e.g., Idle, Active).
Desc	Optional neighbor description (if configured).
Total number of neighbors	Number of configured BGP neighbors.
Total number of Established sessions	Number of neighbors with active, established BGP sessions.

© 2026 IP Infusion Inc. Proprietary

show ip bgp rtfilter neighbors

Use this command to display the BGP Route-Target (RT) filter routes that the router either advertises to or receives from the specified BGP neighbor.

Command Syntax

```
show ip bgp rtfilter neighbors <PREFIX> advertised-routes
show ip bgp rtfilter neighbors <PREFIX> received-routes
```

Parameters

<PREFIX>
Specify the IPv4 address of the BGP neighbor. The command output is specific to this neighbor.

advertised-routes
Displays RT filter routes sent by the local router to the specified neighbor.

received-routes
Displays RT filter routes learned from the specified neighbor.

Command Mode

Execution mode and Privilege mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

Display Advertised-routes

```
OcNOS#show ip bgp rtfilter neighbors 12.1.1.1 advertised-routes
BGP RT-Filter table version is 1, local router ID is 20.1.1.2
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
              l - labeled, S Stale, x-EVPN
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network            Next Hop            Metric    LocPrf   Weight Path
 *> 100:0:200:1        0.0.0.0                                  32768  i
 *> 100:0:65535:1      0.0.0.0                                  32768  i
 *> 100:1:1.1.1.1:100  0.0.0.0                                  32768  i
 *> 100:2:65551:1      0.0.0.0                                  32768  i

 Total number of route-targets 4

 Example: The format 100:2:1:1 is interpreted as ORIGIN_AS:type/subtype:Route Target
          <Origin-AS>     : Autonomous System Number.
          <Type>          : The type/subtype of the extended community attribute.
                            (0-> 2-byte ASN, 1-> IPv4-address, 2-> 4-byte ASN)
          <RT-Identifier> : Route Target itself.
```

Display Received-routes

```
OcNOS#show ip bgp rtfilter neighbors 12.1.1.1 received-routes
BGP RT-Filter table version is 1, local router ID is 20.1.1.2
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
```

```
              l - labeled, S Stale, x-EVPN
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network            Next Hop            Metric    LocPrf   Weight Path
*> 100:0:100:1        0.0.0.0                                 0      i

Total number of route-targets 1

Example: The format 100:2:1:1 is interpreted as ORIGIN_AS:type/subtype:Route Target
         <Origin-AS>     : Autonomous System Number.
         <Type>          : The type/subtype of the extended community attribute.
                           (0-> 2-byte ASN, 1-> IPv4-address, 2-> 4-byte ASN)
         <RT-Identifier> : Route Target itself.
```

Here is the explanation of the show output fields.

Table 41. Description of the output field for the `show ip bgp rtfilter neighbors` command.

Field	Description
BGP RT-Filter table version	Version number of the RT-Filter table. Increments when changes occur.
local router ID	Router ID of the local BGP process. **Example:** 20.1.1.2
Status codes	Symbols indicating route status: `s (suppressed), d (damped), h (history), a (add-path), b (back-up), * (valid), > (best), i (internal), l (labeled), S (stale), x (EVPN).`
Origin codes	Route origin information: `i (IGP), e (EGP), ? (incomplete).`
Network	RT filter in the extended community format `<Origin-AS>:<Type>:<RT-Identifier>`. **Example:** 100:2:1:1 - Origin-AS=100, Type=2 (4-byte ASN), RT-Identifier=1.
Next Hop	IP address of the next hop for the route. For locally originated routes, this is `0.0.0.0`.
Metric	BGP metric for the route (if applicable).
LocPrf	Local preference value applied to the route.
Weight	Local router-specific weight attribute.
Path	AS path associated with the route.
Total number of route-targets	Displays the total count of RT entries advertised or received.

show ip bgp scan

Use this command to display BGP scan status.

Command Syntax

```
show ip bgp scan
```

Parameters

None

Command Mode

Execution modePrivilege mode

Applicability

This command was introduced before OcNOS version 1.3.

In OcNOS version 7.0.0, the `MPLS nexthop: nhlfe-ix` field is added to the `show ip bgp scan` output.

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
```

The `MPLS nexthop: nhlfe-ix` field in the `show ip bgp scan` output indicates that the BGP next-hop is resolved through an MPLS forwarding entry, reflecting integration with the MPLS-aware recursive next-hop resolution mechanism.

```
#show ip bgp scan
    BGP VRF: (Default) VRF_ID 0, AS 200, router-id 1.1.1.1 RIB SCAN FRESH
    BGP scan interval is 60 secs
    scan remain-time: 19 secs
    Current BGP nexthop cache:
     2.2.2.2 valid [IGP metric 2, IfIndex: 10005, LSP status : invalid, LSP Type : NONE, is_dep : No,
MPLS nexthop: nhlfe-ix 1]
```

© 2026 IP Infusion Inc. Proprietary

show ip bgp vpnv4

Use this command to display information relating to VPNv4.

Command Syntax

```
show ip bgp vpnv4 all
show ip bgp vpnv4 all A.B.C.D
show ip bgp vpnv4 all neighbors
show ip bgp vpnv4 all neighbors A.B.C.D
show ip bgp vpnv4 all neighbors A.B.C.D routes
show ip bgp vpnv4 all summary
show ip bgp vpnv4 all tags
show ip bgp vpnv4 rd WORD
show ip bgp vpnv4 rd WORD A.B.C.D
show ip bgp vpnv4 rd WORD label
show ip bgp vpnv4 rd WORD neighbors
show ip bgp vpnv4 rd WORD neighbors A.B.C.D
show ip bgp vpnv4 rd WORD summary
show ip bgp vpnv4 view WORD all
show ip bgp vpnv4 vrf NAME
show ip bgp vpnv4 vrf NAME A.B.C.D
show ip bgp vpnv4 vrf NAME label
show ip bgp vpnv4 vrf NAME summary
```

Parameters

all
Displays information about all VPNv4 NLRIs

A.B.C.D
Network

neighbors
TCP and BGP neighbor connections

A.B.C.D
Network

routes
Display routes learned from neighbor

summary
Summary display

tags
BGP tags for prefixes

rd
Route distinguisher

WORD
BGP view name

A.B.C.D
Network

label
MPLS Labels for prefixes

neighbors
TCP and BGP neighbor connections

© 2026 IP Infusion Inc. Proprietary

A.B.C.D
Network

summary
Summary display

view
VPNv4 NLRI-specific information

WORD
BGP view name

vrf
VRF VPNv4 NLRIs

NAME
VPN Routing/Forwarding instance name

A.B.C.D
Network

label
MPLS Labels for prefixes

summary
Summary display

Command Mode

Execution modePrivilege mode

Applicability

This command was introduced before OcNOS version 1.3. Introduced new parameter `routes` in OcNOS version 6.4.1

Example

This is a sample output from the `show ip bgp vpnv4` command displaying VPNv4 specific information

```
PE2# sh ip bgp vpnv4 rd 100:1
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

 Route Distinguisher: 100:1 (Default for VRF VRF1)
*>i l 2.1.1.0/24       192.168.36.11        0        100       0   65011 i
*>  l 3.1.1.0/24       201.1.1.254          0        100       0   65011 65005 ?
        -
*>  l 10.10.10.10/32   201.1.1.254          0        100       0   65011 ?
*>  l 20.20.20.0/32    201.1.1.254          0        100       0   65011 i
*>i l 101.1.1.0/24     192.168.36.11        0        100       0    ?        -
*>  l 194.0.1.0        201.1.1.254          0        100       0   65011 ?
*>  l 194.0.2.0        201.1.1.254          0        100       0   65011 ?
*>  l 201.1.1.0        0.0.0.0              0        100       32768  ?        -
 VRF Origin Route count = 6
 Accepted routes count = 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
Route Distinguisher: 100:1
*>i l 2.1.1.0/24       192.168.36.11        0        100       0   65011 i
i l                  192.168.36.12        0        100       0   65011 i

*>i l 101.1.1.0/24     192.168.36.11        0        100       0    ?        -

i l                  192.168.36.12        0        100       0    ?        -
VRF Origin Route count = 0
Accepted routes count = 4
PE2#  sh ip bgp vpnv4
all           fast-reroute  flowspec      rd            rr-asbr-hash
view          vrf
PE2#  sh ip bgp vpnv4 all
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
Route Distinguisher: 100:1 (Default for VRF VRF1)
*>i l 2.1.1.0/24       192.168.36.11        0        100       0   65011 i
*>  l 3.1.1.0/24       201.1.1.254          0        100       0   65011 65005 ?
        -
*>  l 10.10.10.10/32   201.1.1.254          0        100       0   65011 ?
*>  l 20.20.20.0/32    201.1.1.254          0        100       0   65011 i
*>i l 101.1.1.0/24     192.168.36.11        0        100       0    ?        -
*>  l 194.0.1.0        201.1.1.254          0        100       0   65011 ?
*>  l 194.0.2.0        201.1.1.254          0        100       0   65011 ?
*>  l 201.1.1.0        0.0.0.0              0        100       32768  ?        -
 VRF Origin Route count = 6
 Accepted routes count = 2
Route Distinguisher: 100:1
*>i l 2.1.1.0/24       192.168.36.11        0        100       0   65011 i
i l                  192.168.36.12        0        100       0   65011 i

*>i l 101.1.1.0/24     192.168.36.11        0        100       0    ?        -

i l                  192.168.36.12        0        100       0    ?        -
VRF Origin Route count = 0
Accepted routes count = 4
Route Distinguisher: 100:2 (Default for VRF VRF2)



3.1.2.0/24       201.2.1.254          0        100       0   65012 i


*>i l 101.1.2.0/24     192.168.36.11        0        100       0    ?        -
*>  l 201.1.2.0        0.0.0.0              0        100       32768  ?        -
 VRF Origin Route count = 2
 Accepted routes count = 1
Route Distinguisher: 100:2
*>i l 101.1.2.0/24     192.168.36.11        0        100       0    ?        -

i l                  192.168.36.12        0        100       0    ?        -
VRF Origin Route count = 0
```

```
Accepted routes count = 2
Route Distinguisher: 192.168.36.1:1001
*>i l 10.0.0.0/24      192.168.36.11        0        100       0    ?        -

i l                  192.168.36.12        0        100       0    ?        -
*>i l 172.168.10.0/24  192.168.36.11        0        100       0   10 i        -

i l                  192.168.36.12        0        100       0   10 i        -
VRF Origin Route count = 0
Accepted routes count = 4
Route Distinguisher: 192.168.36.2:1001 (Default for VRF irb1001)
*>  l 0.0.0.0/0        30.0.0.254           0        100       0   20 ?        -
*>i l 10.0.0.0/24      192.168.36.1         0        100       0    ?        -

i l                  192.168.36.11        0        100       0    ?        -
*>  l 30.0.0.0/24      0.0.0.0              0        100       32768  ?        -
*>i l 30.0.3.0/24      192.168.36.3         0        100       0    ?        -

i l                  192.168.36.11        0        100       0    ?        -
*>i l 30.0.4.0/24      192.168.36.4         0        100       0    ?        -

i l                  192.168.36.11        0        100       0    ?        -
*>i l 172.168.10.0/24  192.168.36.1         0        100       0   10 i        -

i l                  192.168.36.11        0        100       0   10 i        -
*>  l 192.0.1.0        30.0.0.254           0        100       0   20 65005 ?

*>  l 193.0.1.0        30.0.0.254           0        100       0   20 ?        -
 VRF Origin Route count = 4
 Accepted routes count = 8
Route Distinguisher: 192.168.36.3:1001
*>i l 30.0.3.0/24      192.168.36.11        0        100       0    ?        -

i l                  192.168.36.12        0        100       0    ?        -
VRF Origin Route count = 0
Accepted routes count = 2
Route Distinguisher: 192.168.36.4:1001
*>i l 30.0.4.0/24      192.168.36.11        0        100       0    ?        -

i l                  192.168.36.12        0        100       0    ?        -
VRF Origin Route count = 0
Accepted routes count = 2
PE2#
```

Table 42 explains the fields shows for each route.

Table 42. show ip bgp vpnv4 all neighbors output details

Field	Description
BGP neighbor	Router ID of the BGP neighbor.
remote AS	Autonomous system number of the neighbor.
local AS	Autonomous system number of the local system.
internal link	internal link: iBGP neighbor (in the same AS).
BGP version	The version of BGP being used by the neighbor device.
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.

© 2026 IP Infusion Inc. Proprietary

Table 42. show ip bgp vpnv4 all neighbors output details (continued)

Field	Description
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: 1. Route refresh 2. Address family IPv4 Unicast 3. Address family IPv6 Unicast
Received	Message count, notification count, number of messages waiting in the queue.
Sent	Message count, notification count, number of messages waiting in the queue.
Route refresh request	Route requests sent and received.
For address family	As stated – in this case IPv4 Unicast.
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.
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

© 2026 IP Infusion Inc. Proprietary

Table 42. show ip bgp vpnv4 all neighbors output details (continued)

Field	Description
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 ip bgp vpnv6 all neighbors

Use this command to display VPNv6 NLRI information for all neighbors or for a given neighbor.

Command Syntax

```
show ip bgp vpnv6 all neighbors
show ip bgp vpnv6 all neighbors (A.B.C.D|X:X::X:X)
```

Parameters

A.B.C.D
IPv4 neighbor address

X:X::X:X
IPv6 neighbor address

Command Mode

Execution modePrivilege mode

Applicability

This command was introduced in OcNOS version 3.0.

Example

```
#show ip bgp vpnv6 all neighbors
BGP neighbor is 10.32.0.3, remote AS 1, local AS 1, internal link
  BGP version 4, local router ID 10.32.0.2, remote router ID 10.32.0.3
  BGP state = Established, up for 00:37:31
  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
    Address family VPNv4 Unicast: advertised and received
    Address family VPNv6 Unicast: advertised and received
  Received 25711 messages, 0 notifications, 0 in queue
  Sent 25673 messages, 1 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

 For address family: VPNv4 Unicast
  BGP table version 837, neighbor version 837
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: VPNv6 Unicast
  BGP table version 302, neighbor version 302
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
```

© 2026 IP Infusion Inc. Proprietary

```
  0 accepted prefixes
  0 announced prefixes

 Connections established 2; dropped 1
Local host: 10.32.0.2, Local port: 179
Foreign host: 10.32.0.3, Foreign port: 38416
Nexthop: 10.32.0.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:37:34, due to Hold Timer Expired (Notification sent)
Notification Error Message: (Hold Timer Expired/No sub-error code)

BGP neighbor is 101.101.1.2, vrf 1, remote AS 101, local AS 1, external link
  BGP version 4, local router ID 101.101.1.1, remote router ID 0.0.0.0
  BGP state = Idle
  Last read         , hold time is 90, keepalive interval is 30 seconds
  Received 512 messages, 0 notifications, 0 in queue
  Sent 450 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 27, neighbor version 0
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (standard)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 1

BGP neighbor is 101.102.1.2, vrf 2, remote AS 102, local AS 1, external link
  BGP version 4, local router ID 101.102.1.1, remote router ID 0.0.0.0
  BGP state = Idle
  Last read         , hold time is 90, keepalive interval is 30 seconds
  Received 511 messages, 0 notifications, 0 in queue
  Sent 449 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 26, neighbor version 0
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (standard)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 1

BGP neighbor is 101.103.1.2, vrf 3, remote AS 103, local AS 1, external link
  BGP version 4, local router ID 101.103.1.1, remote router ID 0.0.0.0
  BGP state = Idle
  Last read         , hold time is 90, keepalive interval is 30 seconds
  Received 511 messages, 0 notifications, 0 in queue
  Sent 452 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 26, neighbor version 0
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (standard)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 1

BGP neighbor is 101.104.1.2, vrf 4, remote AS 104, local AS 1, external link
  BGP version 4, local router ID 101.104.1.1, remote router ID 0.0.0.0
  BGP state = Connect
  Last read         , hold time is 90, keepalive interval is 30 seconds
```

```
  Received 511 messages, 0 notifications, 0 in queue
  Sent 449 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 27, neighbor version 0
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (standard)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 1
Next connect timer due in 108 seconds

BGP neighbor is 101.105.1.2, vrf 5, remote AS 105, local AS 1, external link
  BGP version 4, local router ID 101.105.1.1, remote router ID 0.0.0.0
  BGP state = Idle
  Last read         , hold time is 90, keepalive interval is 30 seconds
  Received 511 messages, 0 notifications, 0 in queue
  Sent 454 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 26, neighbor version 0
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (standard)
  0 accepted prefixes
  0 announced prefixes

 Connections established 1; dropped 1
```

show ip bgp vpnv6 rd neighbors

Use this command to display VPNv6 NLRI information for a given route distinguisher and optionally for a given neighbor.

Command Syntax

```
show ip bgp vpnv6 rd WORD neighbors
show ip bgp vpnv6 rd WORD neighbors (A.B.C.D|X:X::X:X)
```

Parameters

WORD
Route distinguisher

A.B.C.D
IPv4 neighbor address

X:X::X:X
IPv6 neighbor address

Command Mode

Execution modePrivilege mode

Applicability

This command was introduced in OcNOS version 3.0.

Example

```
#show ip bgp vpnv6 rd 1:1 neighbors
BGP neighbor is 10.32.0.3, remote AS 1, local AS 1, internal link
  BGP version 4, local router ID 10.32.0.2, remote router ID 10.32.0.3
  BGP state = Established, up for 00:43:57
  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
    Address family VPNv4 Unicast: advertised and received
    Address family VPNv6 Unicast: advertised and received
  Received 25726 messages, 0 notifications, 0 in queue
  Sent 25689 messages, 1 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

 For address family: VPNv4 Unicast
  BGP table version 837, neighbor version 837
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
```

```
 0 announced prefixes

For address family: VPNv6 Unicast
 BGP table version 302, neighbor version 302
 Index 1, Offset 0, Mask 0x2
 Community attribute sent to this neighbor (both)
 0 accepted prefixes
 0 announced prefixes

Connections established 2; dropped 1
```

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

Execution modePrivilege 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
```

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivilege 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 43 explains the output fields.

Table 43. 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.

© 2026 IP Infusion Inc. Proprietary

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

Execution modeand Privilege 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
```

© 2026 IP Infusion Inc. Proprietary

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#
```

© 2026 IP Infusion Inc. Proprietary

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)#
```

© 2026 IP Infusion Inc. Proprietary

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#
```

© 2026 IP Infusion Inc. Proprietary

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

Configure mode

Applicability

This command was introduced in OcNOS version 6.6.0.

Example

```
OcNOS#show running-config ip vrf
ip vrf management
!
ip vrf test2
!
ip vrf test5
!
!
end
OcNOS#
```

© 2026 IP Infusion Inc. Proprietary

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

Configure mode

Applicability

This command was introduced in OcNOS version 6.6.0.

Example

```
OcNOS#show running-config mac vrf
mac vrf test1
!
mac vrf test3
!
mac vrf test4
!
!
end
OcNOS#
```

© 2026 IP Infusion Inc. Proprietary

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#
```

BGP Additional Paths Commands

The section describes the BGP additional paths commands.

Notes:

- The following protocols support additional paths: VPNv4, VPNv6, 6PE, and IPv4 Unicast (without IPv4 Labeled Unicast).
- Additional path advertisement is only supported for Internal Border Gateway Protocol (IBGP) peers. Furthermore, the marking for additional paths is performed for internal paths.
- When the next-hop-self configuration is set for a BGP peer, the device will only advertise the best path to add-path neighbors where the next-hop-self is enabled.
- A maximum of 64 paths can be advertised for the same prefix using the additional path.
- The system assigns a permanent, immutable tx path_id to every route upon creation (supporting Unicast, Labeled Unicast, and VPN SAFIs), replacing previous default values (e.g., 65 or -1) with stable 32-bit integer tracking.

bgp additional-paths 1595
bgp additional-paths install 1596
neighbor additional-paths 1597
neighbor advertise additional-paths 1598

© 2026 IP Infusion Inc. Proprietary

bgp additional-paths

Use this command to enable BGP additional paths.

Use the `no` parameter with this command to disable BGP add-path.

Command Syntax

```
bgp additional-paths (send|receive|send-receive|select (all | best<2-3>))
no bgp additional-paths (send|receive|send-receive|select (all | best))
```

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

all
Select all available paths

best
Select best N paths

<2-3>
Number of best paths in additional paths to be selected

Default

Disabled

Command Mode

address-family vpnv4 unicast

address-family vpnv6 unicast

address-family ipv6 labeled-unicast

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 2
(config-router)#address-family ipv4 unicast
(config-router-af)#bgp additional-paths send
(config-router-af)#no bgp additional-paths send
```

© 2026 IP Infusion Inc. Proprietary

bgp additional-paths install

Use this command to install the back-up path. BGP selects only one best path in the multiple next-hop network, the additional-path install (or) bgp prefix independent convergence enables BGP to select and install the back-up path.

Use the no parameter with this command to delete the installed back-up paths.

Command Syntax

```
bgp additional-paths install
no bgp additional-paths install
```

Parameters

None

Default

Disabled

Command Mode

address-family vpnv4 unicast

address-family vpnv6 unicast

address-family ipv6 labeled-unicast

Applicability

This command was introduced before OcNOS version 4.1 and new modes were introduced before OcNOS version 6.2.0

Examples

```
OCNOS(config)#router bgp 100
OCNOS(config-router)#address-family vpnv4 unicast
OCNOS(config-router)#bgp additional-paths install
OCNOS(config-router-af)#exit-address-family

OCNOS(config-router)#address-family vpnv6 unicast
OCNOS(config-router)#bgp additional-paths install
OCNOS(config-router-af)#exit-address-family

OCNOS(config-router)#address-family ipv6 labeled-unicast
OCNOS(config-router)#bgp additional-paths install
OCNOS(config-router-af)#exit-address-family
```

CLI ”bgp additional-paths select best” requires to be configured.

© 2026 IP Infusion Inc. Proprietary

neighbor additional-paths

Use this command to enable neighbour BGP add-path.

Use the `no` parameter with this command to disable neighbour BGP add-path.

Command Syntax

```
neighbor A.B.C.D additional-paths (send|receive|send-receive|disable)
no neighbor A.B.C.D additional-paths (send|receive|send-receive|disable)
```

Parameters

send
Send additional paths to neighbors

receive
Receive additional paths from neighbors

send-receive
Send and Receive additional paths from neighbors

disable
Disable additional paths

Default

Disabled

Command Mode

Address Family mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 2
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 1.1.1.2 additional-paths send
(config-router-af)#no neighbor 1.1.1.2 additional-paths send
```

© 2026 IP Infusion Inc. Proprietary

neighbor advertise additional-paths

Use this command to enable BGP add-path at neighbor level.

Use the `no` parameter with this command to disable BGP add-path at neighbor level.

Command Syntax

```
neighbor A.B.C.D advertise additional-paths (all|best <2-3>)
no neighbor A.B.C.D advertise additional-paths (all|best <2-3>)
```

Parameters

all

Select all available paths

best

Select best N paths

<2-3>

Number of best paths in additional paths to be selected

Default

Disabled

Command Mode

Address Family mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router bgp 2
(config-router)#address-family ipv4 unicast
(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
```

© 2026 IP Infusion Inc. Proprietary

BGP Graceful Restart Commands

This section describes the BGP graceful restart commands.

bgp g-shut 1600
bgp g-shut-capable 1601
bgp g-shut-local-preference 1602
bgp graceful-restart 1603
bgp update-delay 1605
neighbor capability graceful-restart 1606
neighbor g-shut 1608
neighbor g-shut-timer 1609
neighbor restart-time 1610

© 2026 IP Infusion Inc. Proprietary

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 the maintenance activity using the `bgp g-shut` command.

Note: The graceful shutdown is not supported on iBGP sessions.

Command Syntax

```
bgp g-shut
no bgp g-shut
```

Parameters

None

Default

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
```

© 2026 IP Infusion Inc. Proprietary

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.

Notes:

- 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.
- The graceful shutdown capability is not supported on iBGP sessions.

Command Syntax

```
bgp g-shut-capable
no bgp g-shut-capable
```

Parameters

None

Default

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 (page 1602) command relies on the bgp g-shut-capable (page 1601) configuration. When user remove the bgp g-shut-capable (page 1601) command from the configuration, it automatically deletes the bgp g-shut-local-preference (page 1602) setting as well. This dependency means that the bgp g-shut-local-preference (page 1602) cannot function or remain configured without the presence of the bgp g-shut-capable (page 1601) command.

Command Syntax

```
bgp g-shut-local-preference <0-4294967295>
no bgp g-shut-local-preference
```

Parameters

<0-4294967295>
Local preference value

Default

Zero

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 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>
sDelay 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

The maximum time for neighbors to restart is 90 seconds.

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
```

© 2026 IP Infusion Inc. Proprietary

```
(config)#router bgp 10
(config-router)#bgp graceful-restart

#configure terminal
(config)#router bgp 10
(config-router)#no bgp graceful-restart
```

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

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1603) 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 (page 1411) command. When you specify this parameter, the command applies to all peers in the group.

Default

Disabled

Command Mode

Address Family mode Applicable for ipv4, ipv6 and vpnv4.

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
```

```
(config-router)#address-family vpnv4 unicast
(config-router-af)#neighbor 2.2.2.2 activate
(config-router-af)#neighbor 2.2.2.2 capability graceful-restart
(config-router-af)#neighbor 3.3.3.3 activate
(config-router-af)#neighbor 3.3.3.3 capability graceful-restart
```

© 2026 IP Infusion Inc. Proprietary

```
(config-router-af)#exit-address-family
```

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 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 (page 1603) 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 (page 1411) 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

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
```

© 2026 IP Infusion Inc. Proprietary

BGP Labeled Unicast Commands

This section describes the BGP labeled unicast commands.

address-family labeled-unicast 1612
allocate-label 1614
clear ip bgp * ipv4 labeled-unicast 1615
clear ip bgp A.B.C.D ipv4 labeled-unicast 1616
clear ip bgp peer-group WORD ipv4 labeled-unicast 1617

© 2026 IP Infusion Inc. Proprietary

address-family labeled-unicast

Use the address family command to enter the IPv4 or VPNv4 address family mode allowing configuration of address-family specific parameters. To leave address family mode and return to configure mode, give the exit-address-family (page 1319) 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.

Use the `no` parameter with this command to disable the address-family configurations.

As in OcNOS, IPv4 Unicast and IPv4 Labeled-Unicast share the same BGP RIB [AFI=1][SAFI=1].

Paths learned from IPv4 unicast capable peer and routes learned from IPv4 labeled-unicast capable peer both the paths will be imported to IPv4 Unicast RIB. For a BGP route A.B.C.D, paths learned via IPv4 Labeled-Unicast are preferred over paths learned via IPv4 Unicast.

When max-paths ECMP is enabled for IPv4 Unicast address-family routes:

For a BGP route having the paths learned via IPv4 Labeled-Unicast and also via IPv4 Unicast OcNOS will not select these paths as multipath candidates as load-balancing the traffic to ip path and mpls path is not achievable in OcNOS and paths from LU are preferred and are marked best path.

Note: Enabling the `address-family ipv4 labeled-unicast` mode, automatically adds bgp implicit-null CLI to the existing configuration. For more information, refer to .

Command Syntax

```
address-family ipv4 labeled-unicast
address-family ipv6 labeled-unicast
```

Parameters

ipv4

IPv4 address family

labeled-unicast

Enter IPv4 labeled-unicast address-family mode to advertise labeled unicast routes. When a neighbor activate (page 1338) command is given in this mode, the BGP speaker advertises the BGP-LU capability.

ipv6

IPv6 address family

labeled-unicast

Enter IPv6 labeled-unicast address-family mode to:

Activate an IPv4 neighbor to exchange labeled routes data among ISP PE devices. When a IPv4 neighbor activate (page 1338) command is given in this mode, the device becomes 6PE capable.

OR:

Activate an IPv6 neighbor to advertise labeled unicast routes. When a IPv6 neighbor activate (page 1338) command is given in this mode, the BGP speaker advertises the BGP-LU capability.

© 2026 IP Infusion Inc. Proprietary

Default

None

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3 and updated in OcNOS version 3.0.

Example

```
(config)#router bgp 100
(config-router)#address-family ipv4 labeled unicast
(config-router-af)#neighbor 172.4.5.52 activate
(config-router-af)#exit-address-family
```

© 2026 IP Infusion Inc. Proprietary

allocate-label

Use this command to allocate MPLS labels to all/route-map permitted IPv4 prefixes present in the BGP RIB. This command enables the label allocation to IPv4 prefixes, where labeled unicast (LU) routes can be exchanged with IPv4 Peers having BGP labeled-unicast capability. This command allocates labels for all IPv4 unicast prefixes (afi:1, safi:1) and IPv4 labeled-unicast prefixes (afi:1, safi:4).

Once configured, BGP will allocate labels to IPv4 unicast prefixes after the expiry of current BGP RIB scan time (default: 60 seconds).

Use the `no` form of this command to release the labels and disable BGP from allocating label to IPv4 prefixes.

Command Syntax

```
allocate-label (all|route-map WORD)
no allocate-label (all|route-map)
```

Parameters

all
All the advertised routes

route-map
Only those routes that match route-map

WORD
Route-map name

Default

None

Command Mode

Router mode

Applicability

This command was introduced before updated in OcNOS version 3.0.

Example

To allocate an MPLS label to IPv4 unicast prefix 1.1.1.1/32

```
(config)#router bgp 100
(config-router)#neighbor 11.0.0.2 remote-as 100
(config-router)#allocate-label all
(config-router)#address-family ipv4 unicast
(config-router-af)#network 1.1.1.1/32
(config-router-af)#exit-address-family
(config-router)#address-family ipv4 labeled-unicast
(config-router-af)#network 11.0.0.2 activate
(config-router-af)#exit-address-family
```

```
(config-router)#no allocate-label all
(config-router)#allocate-label route-map bgp
(config-router)#no allocate-label route-map
(config-router)#
```

© 2026 IP Infusion Inc. Proprietary

clear ip bgp * ipv4 labeled-unicast

Use this command to soft reset all IPv4 Labeled-unicast BGP neighbors.

Command Syntax

```
clear ip bgp * ipv4 labeled-unicast soft
clear ip bgp * ipv4 labeled-unicast soft in
clear ip bgp * ipv4 labeled-unicast soft out
```

Parameters

soft
Soft clear both incoming and outgoing routes

in
Soft reconfig inbound update

out
Soft reconfig outbound update

Command Mode

Privileged execution mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#clear ip bgp * ipv4 labeled-unicast soft in
```

© 2026 IP Infusion Inc. Proprietary

clear ip bgp A.B.C.D ipv4 labeled-unicast

Use this command to soft reset an IPv4 Labeled-unicast BGP neighbor.

Command Syntax

```
clear ip bgp A.B.C.D ipv4 labeled-unicast soft
clear ip bgp A.B.C.D ipv4 labeled-unicast soft in
clear ip bgp A.B.C.D ipv4 labeled-unicast soft out
```

Parameters

A.B.C.D
BGP neighbor address

soft
Soft clear both incoming and outgoing routes

in
Soft reconfig inbound update

out
Soft reconfig outbound update

Command Mode

Privileged execution mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#clear ip bgp A.B.C.D ipv4 labeled-unicast soft in
```

© 2026 IP Infusion Inc. Proprietary

clear ip bgp peer-group WORD ipv4 labeled-unicast

Use this command to soft reset an peer-group of ipv4 labeled-unicast peer members.

Command Syntax

```
clear ip bgp peer-group WORD ipv4 labeled-unicast soft
clear ip bgp peer-group WORD ipv4 labeled-unicast soft in
clear ip bgp peer-group WORD ipv4 labeled-unicast soft out
```

Parameters

WORD
Peer-group name

soft
Soft clear both incoming and outgoing routes

in
Soft reconfig inbound update

out
Soft reconfig outbound update

Command Mode

Privileged execution mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#clear ip bgp peer-group LU_GROUP ipv4 labeled-unicast soft in
```

© 2026 IP Infusion Inc. Proprietary

Regular Expressions

The table below shows the regular expression special characters used in BGP commands. Users can use these characters in combination to build complex regular expressions.

Table 44. 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.

© 2026 IP Infusion Inc. Proprietary

BIDIRECTIONAL FORWARDING DETECTION CONFIGURATION

BFD Authentication 1621
Overview 1621
Enabling BFD Authentication for Multihop for IPv4, Using Key-ID 1621
Topology 1621
Configuration 1622
Base BFD Configuration 1626
Overview 1626
BFD Echo Function 1627
BFD Slow Timer 1629
BFD Multihop Peer Timer 1630
BFD Single-hop Session Timer 1633
BFD Echo Interval 1635
BFD Protocol Configurations 1637
OSPF—BFD Single-Hop Session 1637
OSPF—BFD Multi-Hop Session 1644
BFD Configuration in IS-IS 1647
BFD Configuration in BGP 1649
BFD Static Route Configuration 1657
Topology 1657
Validation 1658
BFD with VRF Configuration 1661
Topology 1661
RTR1 Configuration 1661
Switch Configuration 1662
RTR2 Configuration 1663
BFD with VRF Validation 1664
BFD Over Static Routing IPv4 and IPv6 1667
BFD Support on LAG Interface 1671
Overview 1671
Single Hop BFD Option Over LAG 1671
Feature Characteristics 1671
Benefits 1671
Prerequisites 1672
Configuration 1672
Validation 1673

© 2026 IP Infusion Inc. Proprietary

CLI Commands 1674

© 2026 IP Infusion Inc. Proprietary

BFD Authentication

This section 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 101. Basic Topology of Three Routers

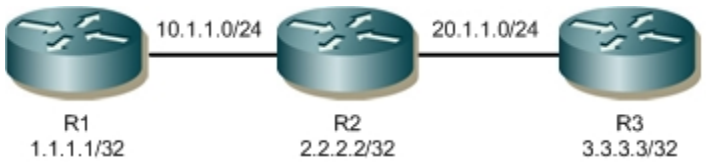

© 2026 IP Infusion Inc. Proprietary

Configuration

R1 Configuration

`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

R2 Configuration

`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

© 2026 IP Infusion Inc. Proprietary

`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

R3 Configuration

`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 : 00000000000000000000            IPv6 Echo Out : 00000000000000000000
IPv6 Pkt In : 00000000000000000000         IPv6 Pkt Out : 00000000000000000000
UP Count : 12                              UPTIME : 00:36:29

Protocol Client Info:
BFD-> Client ID: 28     Flags: 4
------------------------------------------------------------
Number of Sessions:    1
```

Base BFD Configuration

This section 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 sections in this section.

Note: Software BFD is enabled by default on IRB interfaces, as hardware BFD is not supported.

Topology

Figure 102. Basic Topology of Three Routers

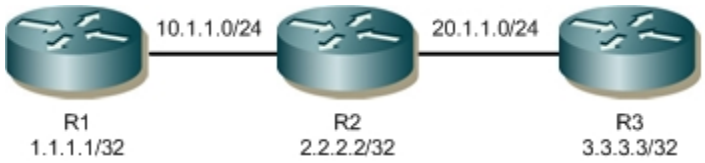

BFD Echo Function

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

```
#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 : 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

```
#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 : 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
```

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.

© 2026 IP Infusion Inc. Proprietary

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

```
#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 : 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
```

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 : 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
```

R2

```
#show 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
```

© 2026 IP Infusion Inc. Proprietary

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

```
#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 : 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

```
#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 : 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
```

BFD Protocol Configurations

This section describes the BFD protocol configurations.

OSPF—BFD Single-Hop Session

This section provides the steps for configuring BFD for Single-Hop OSPF.

Topology

Figure 103. Single-Hop OSPF Topology

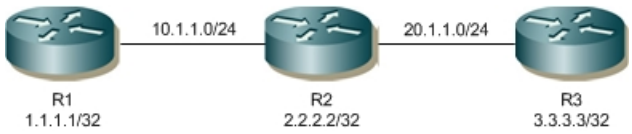

R1

`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 `area1`.
`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

© 2026 IP Infusion Inc. Proprietary

`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 104. Multi-hop OSPFv2 Topology

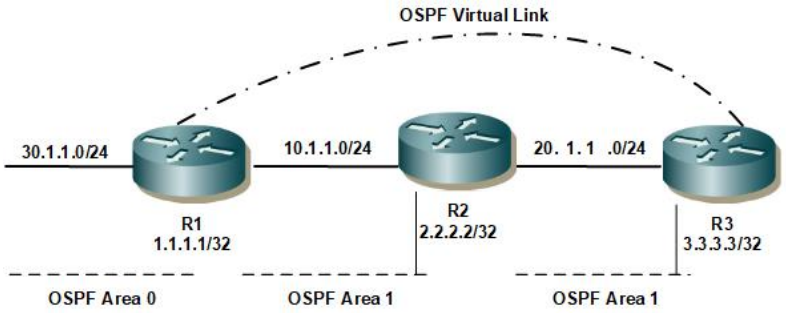

R1

`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`	
`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.
`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 105. Basic Topology for BFD-ISIS

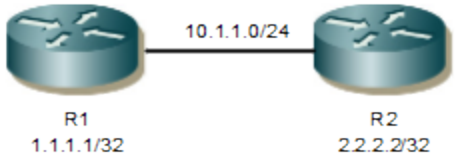

R1

`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 106. Basic Topology for BFD in BGP

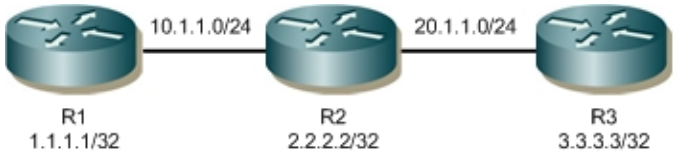

R1

`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

© 2026 IP Infusion Inc. Proprietary

`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.

© 2026 IP Infusion Inc. Proprietary

`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
```

BFD Static Route Configuration

This section 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 107. BFD Static Route Basic Topology

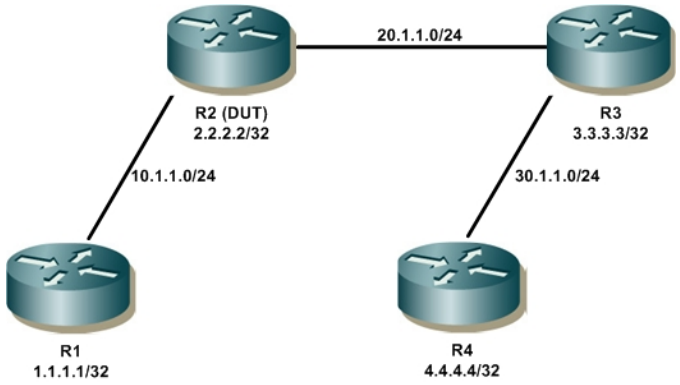

R1

`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.

© 2026 IP Infusion Inc. Proprietary

`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
```

BFD with VRF Configuration

This chapter shows using BFD with user defined VRF for OSPFv2 and OSPFv3.

Topology

Figure 108. BFD user-defined VRF

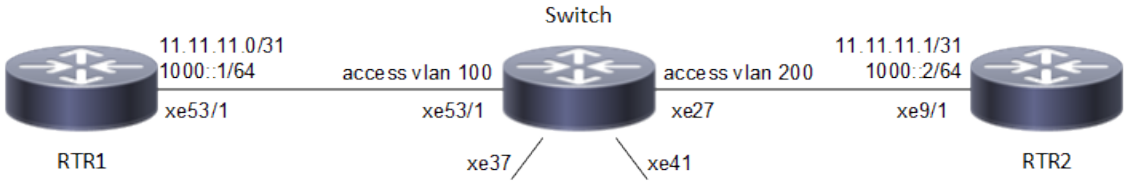

RTR1 Configuration

`#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

© 2026 IP Infusion Inc. Proprietary

(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.

Switch Configuration

#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

© 2026 IP Infusion Inc. Proprietary

(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 Configuration

#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.

© 2026 IP Infusion Inc. Proprietary

BFD with VRF 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 Configuration

#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 Configuration

#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

© 2026 IP Infusion Inc. Proprietary

(config)#interface xe9/1	Enter interface mode f or xe9/1
(config-if)#ip vrf forwarding vrf10	Enable VRF forwarding on interface
(config-if)#ip address 11.11.11.1/31	A s si gn IPv4 ad dr e s s in /31 mask.
(config-if)#ipv6 address 1000::2/64	A s si gn IPv6 ad dr e s s 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

BFD Over Static Routing Validation

RTR1

```
#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
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
```

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.

Notes:

- When setting up the session with LAG using the " bfd session software" command, the session type is single-hop.
- When setting up the session with LAG without using the command " bfd session software" command, the session type is 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).

© 2026 IP Infusion Inc. Proprietary

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 and R2.

Figure 109. BFD support on LAG interface

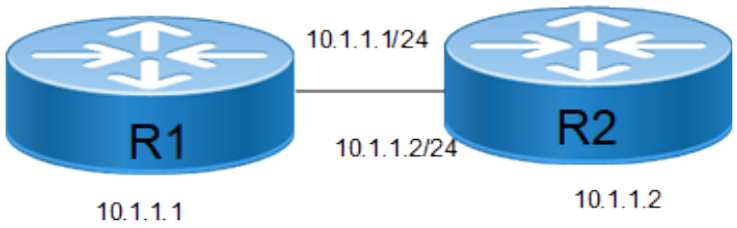

Perform the following steps to configure the BFD on the LAG interface.

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
```

CLI Commands

bfd session software

Use this command to create a BFD session operated by control plane (This implies the BFD TX or RX is 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

None

Command Mode

Interface mode

Applicability

Introduced in 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.

© 2026 IP Infusion Inc. Proprietary

BIDIRECTIONAL FORWARDING DETECTION COMMAND REFERENCE

Bidirectional Forwarding Commands 1678
accept-lifetime 1679
bfd auth type 1681
bfd 1683
bfd echo 1684
bfd echo interval 1685
bfd echo ipv4 source 1686
bfd interval 1687
bfd multihop-peer 1688
bfd multihop-peer A.B.C.D interval 1690
bfd multihop-peer interval 1691
bfd multihop-peer X:X::X:X interval 1692
bfd notification 1693
bfd session 1694
bfd slow-timer 1696
debug bfd 1697
hardware-profile seamless-bfd 1698
key-id 1699
key chain 1700
key-string 1701
key-string encrypted 1702
send-lifetime 1703
show bfd 1705
show bfd interface 1706
show bfd session 1708
show bfd session A.B.C.D 1713
show bfd session ipv6 1716
show debugging bfd 1719
snmp restart bfd 1720
Protocol Commands for BFD 1721
area virtual-link 1722
bfd all-interfaces 1723
debug bgp bfd 1724
debug isis bfd 1725
debug ospf bfd 1726

© 2026 IP Infusion Inc. Proprietary

ip ospf bfd 1727
isis bfd 1728
BFD Static Route Commands 1729
ip bfd static all-interfaces 1730
ip static fall-over-bfd 1731
ip static bfd 1732
ipv6 bfd static all-interfaces 1733
ipv6 static fall-over-bfd 1734
ipv6 static bfd 1735

© 2026 IP Infusion Inc. Proprietary

Bidirectional Forwarding Commands

This chapter describes the Bidirectional Forwarding commands:

accept-lifetime 1679
bfd auth type 1681
bfd 1683
bfd echo 1684
bfd echo interval 1685
bfd echo ipv4 source 1686
bfd interval 1687
bfd multihop-peer 1688
bfd multihop-peer A.B.C.D interval 1690
bfd multihop-peer interval 1691
bfd multihop-peer X:X::X:X interval 1692
bfd notification 1693
bfd session 1694
bfd slow-timer 1696
debug bfd 1697
hardware-profile seamless-bfd 1698
key-id 1699
key chain 1700
key-string 1701
key-string encrypted 1702
send-lifetime 1703
show bfd 1705
show bfd interface 1706
show bfd session 1708
show bfd session A.B.C.D 1713
show bfd session ipv6 1716
show debugging bfd 1719
snmp restart bfd 1720

© 2026 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 disable it.

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 expire.

duration
Specify the duration of the key in seconds

<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

Key-chain key mode

© 2026 IP Infusion Inc. Proprietary

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the setting of accept-lifetime for key-id 1 on the key chain named mychain.

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key-id 1
(config-keychain-key)#accept-lifetime 03:03:01 Dec 03 2004 04:04:02 Oct 06 2006
```

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 off.

Note: BFD echo mode is supported for IPv4 BFD single hop sessions only.

Command Syntax

```
bfd echo
no bfd echo
```

Parameters

None

Default

By default, `bfd echo` is off.

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.

Default

The default value is 1000 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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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.

Notes: 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.

The scaling number of BFD sessions depends on the target device's CPU performance, as 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 interval of 100ms (default interval is 250ms). The lowest configurable interval for a software-based BFD session is 50ms.

Command Syntax

```
bfd interval <3-999> minrx <3-999> multiplier <3-50>
no bfd interval
```

Parameters

<3-999>
Transmit interval in milliseconds.

minrx <3-999>
Receive interval in milliseconds.

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)#
```

© 2026 IP Infusion Inc. Proprietary

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 an IPv6 address.

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

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 the control plane. When handling numerous services or traffic sessions, BFD session instability might occur. To address this, it is recommended to increase 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 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

```
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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

```
bfd session (hardware| Software)(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
Specifies the source IPv4 address for the BFD session.

X:X::X:X
Specifies the source IPv6 address for the BFD session.

hardware
Offloads the BFD session to the forwarding plane (hardware-based processing).

software
Runs the BFD session on the control plane (handled by the host CPU).

multihop
Enables BFD session support over multihop paths

admin-down
Brings the BFD session administratively down.

demand-mode
Enables demand mode for the BFD session, reducing control message exchanges.

non-persistent
Configures the BFD session as non-persistent; session state is not retained across restarts.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3. Added parameters hardware and software in OcNOS version 6.6.1

Examples

```
#configure terminal
(config)#interface eth2
(config-if)#bfd session 20.20.20.2 20.20.20.1 admin-down
(config-if)#do sh run bfd
!
interface eth2
bfd session  20.20.20.2  20.20.20.1 admin-down
!
(config-if)#no bfd session  20.20.20.2  20.20.20.1 admin-down
(config-if)#do sh run bfd
```

© 2026 IP Infusion Inc. Proprietary

```
!
interface eth2
bfd session  20.20.20.2  20.20.20.1
!
(config-if)#bfd session 20.20.20.2 20.20.20.1 admin-down
(config-if)#do sh run bfd
!
interface eth2
bfd session  20.20.20.2  20.20.20.1 admin-down
!
(config-if)#no bfd session 20.20.20.2 20.20.20.1
(config-if)#do sh run bfd
!
```

© 2026 IP Infusion Inc. Proprietary

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 <1000-1703>
no bfd slow-timer
```

Parameter

<1000-1703>
Interval for the slow-timer in milliseconds.

Command Mode

Configure mode

Default

The default slow timer interval is 1703 milli-seconds.

The default slow timer interval is 2000 milli-seconds.

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#bfd slow-timer 1500
(config)#
```

© 2026 IP Infusion Inc. Proprietary

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

Configure mode, Privileged execution mode and Execution 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
```

© 2026 IP Infusion Inc. Proprietary

hardware-profile seamless-bfd

Use this command to enable the hardware-profile for seamless BFD to successfully activate seamless BFD in the hardware. Configure the `hardware-profile seamless-bfd` command before configuring the `s-bfd sr policy` or `s-bfd discriminator` commands. Before disabling the hardware-profile, unconfigure the **s-bfd sr policy** or `s-bfd discriminator` commands.

Use the `disable` form of this command to disable the configured hardware-profile.

Notes:

- After configuring `hardware-profile seamless-bfd` command, save the configuration and reboot the system.
- When downgrading to version 6.3.x from a higher-build version, configure the hardware-profile seamless-bfd disable.

Command Syntax

```
hardware-profile seamless-bfd enable
hardware-profile seamless-bfd disable
```

Parameter

None

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.1.0.

Example

```
#configure terminal
(config)#hardware-profile seamless-bfd enable
(config)#hardware-profile seamless-bfd disable
```

© 2026 IP Infusion Inc. Proprietary

key-id

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-id <0-2147483647>
no key-id <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

Key-chain 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-id 1
(config-keychain-key)#
```

© 2026 IP Infusion Inc. Proprietary

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

Key-chain mode and Key-chain key mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

In the following example, the password for `key-id 1` in the key chain named `mychain` is set to `prime`:

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key-id 1
(config-keychain-key)#key-string prime

(config-keychain)#key-id 1
(config-keychain-key)#no key-string
```

© 2026 IP Infusion Inc. Proprietary

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

Key-chain mode and Key-chain key mode

Applicability

This command was introduced in OcNOS version 4.1.

Example

In the following example, the encrypted password for key-id 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-id 1
(config-keychain-key)#key-string encrypted 0xd6c50b442de47f70
(config-keychain)#key-id 1
(config-keychain-key)#no key-string
```

© 2026 IP Infusion Inc. Proprietary

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

Key-chain key mode

Applicability

This command was introduced before OcNOS version 1.3.

© 2026 IP Infusion Inc. Proprietary

Examples

The following example shows the setting of `send-lifetime` for `key-id 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
```

© 2026 IP Infusion Inc. Proprietary

show bfd

Use this command to display information about the BFD process.

Command Syntax

```
show bfd
```

Parameters

None

Command Mode

Execution mode and Privileged execution 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 45. 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.

© 2026 IP Infusion Inc. Proprietary

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

Privileged execution mode and Execution 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 46 explains the output fields.

© 2026 IP Infusion Inc. Proprietary

Table 46. 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.

© 2026 IP Infusion Inc. Proprietary

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

Privileged execution mode and Execution 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 47 explains the output fields.

Table 47. 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 MPLS LSP MPLS VCCV MPLS-TP
Version	Session version number; generally 1.

Table 47. show bfd session detail fields (continued)

Entry	Description
Session Type Sess-Type	Single Hop Multihop Arbit Path Multihop OOB Signalled Multihop Unidirectional.
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

© 2026 IP Infusion Inc. Proprietary

Table 47. show bfd session detail fields (continued)

Entry	Description
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.
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.

© 2026 IP Infusion Inc. Proprietary

Table 47. show bfd session detail fields (continued)

Entry	Description
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

Privileged execution mode and Execution 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
```

© 2026 IP Infusion Inc. Proprietary

```
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 48 explains the output fields.

Table 48. 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.
Min Rx	Minimum receive interval in milliseconds.
Multiplier	The negotiated transmit interval, multiplied by this value, provides the detection time for

© 2026 IP Infusion Inc. Proprietary

Table 48. Show BFD session fields (continued)

Entry	Description
	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

Privileged execution mode and Execution 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 49 explains the output fields.

Table 49. 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.
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.

© 2026 IP Infusion Inc. Proprietary

Table 49. show BFD session fields (continued)

Entry	Description
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

© 2026 IP Infusion Inc. Proprietary

show debugging bfd

Use this command to display debugging information for BFD processes.

Command Syntax

```
show debugging bfd
```

Parameters

None

Command Mode

Privileged execution mode and Execution 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
#
```

© 2026 IP Infusion Inc. Proprietary

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.

Example

```
#snmp restart bfd
```

© 2026 IP Infusion Inc. Proprietary

Protocol Commands for BFD

The section describes the commands used to manage BFD functionality for OSPF, IS-IS and BGP.

area virtual-link 1722
bfd all-interfaces 1723
debug bgp bfd 1724
debug isis bfd 1725
debug ospf bfd 1726
ip ospf bfd 1727
isis bfd 1728

© 2026 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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#debug bgp bfd
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#debug isis bfd
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#debug ospf bfd
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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.

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
```

© 2026 IP Infusion Inc. Proprietary

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 section includes the following commands:

ip bfd static all-interfaces 1730
ip static fall-over-bfd 1731
ip static bfd 1732
ipv6 bfd static all-interfaces 1733
ipv6 static fall-over-bfd 1734
ipv6 static bfd 1735

© 2026 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
```

© 2026 IP Infusion Inc. Proprietary

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.

Note: Configuring static BFD with an invalid next-hop requires recursive resolution to form a session. However, the system does not support multi-hop static BFD.

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.

fall-over-bfd disable
Disable BFD.

fall-over-bfd enable
Enable BFD.

vrf 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. Added a note in OcNOS version 6.6.1.

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
```

© 2026 IP Infusion Inc. Proprietary

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.

Note: When static bfd is enabled on the interface (`ip static bfd enable`) and if the bfd session is not up, then the related static route to that interface will fall into inactive state.

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

INTERMEDIATE SYSTEM TO INTERMEDIATE SYSTEM CONFIGURATION

IS-IS IPv4 1738
Enable IS-IS on an Interface Configuration 1738
Set Priority 1741
Dynamic Hostname 1743
Redistribute Routes into IS-IS 1745
Metric 1750
L1-L2 Area Routing with a Single Instance 1758
L1-L2 Area Routing with Multiple Instances 1764
Route Leaking 1769
Route Summarization 1774
IS-IS Distance 1779
Overload Bit 1785
Passive Interface 1805
IS-IS IPv4 Loop-Free Alternate Fast Reroute 1810
Validation 1831
IS-IS IPv6 Configuration 1833
Enable IS-ISv6 on an Interface 1833
Set Priority 1835
Dynamic hostname 1837
Redistribute Routes into IS-IS 1840
Interface Metric 1842
Route Summarization 1847
Passive Interface 1852
Enable BFD over IS-ISv6 1857
Originate Default Route to ISISv6 Neighbors 1859
IS-IS Distance 1862
IS-IS IPv4 TE-Metric Extension 1865
Overview 1865
Configuration 1865
Validation 1866
IS-IS Graceful Restart Configuration 1870
Overview 1870
Configuration 1870
IS-IS-TE IPv4 1873
Enable MPLS-TE in Level-1 and Level-2 on L1-L2 IS 1873

© 2026 IP Infusion Inc. Proprietary

Maximum Link Bandwidth and Reservable Bandwidth 1879
Administrative Group Constraints 1887
ISIS Multi Topology 1895
Overview 1895
Prerequisites 1895
Configuration 1898
Validation for Multi Topology 1899
CLI Commands 1923
Glossary 1924

IS-IS IPv4

This section contains basic Intermediate System to Intermediate System (IS-IS) configuration examples.

Enable IS-IS on an Interface Configuration

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 110. Basic IS-IS Topology

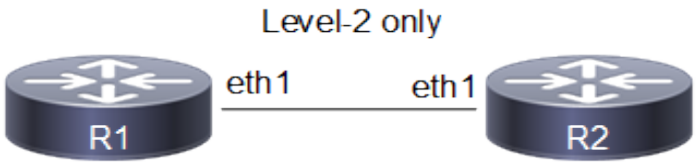

R1 Configuration

`#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.
`(config-router)#commit`	Commit candidate configuration to the running configuration

R2 Configuration

`#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).`

© 2026 IP Infusion Inc. Proprietary

`(config-if)#commit`	Commit candidate configuration to the running configuration
`(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 111. Set IS-IS Priority

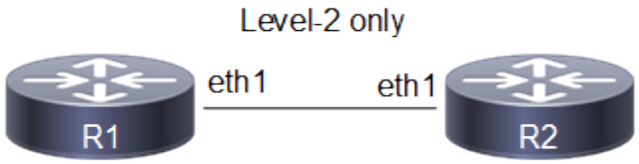

R1 Configuration

`(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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration

R2 Configuration

`(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)#commit`	Commit candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area `49` (`ABC`).

© 2026 IP Infusion Inc. Proprietary

`(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 112. Basic dynamic hostname topology

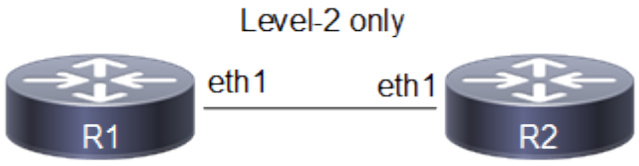

R1 Configuration

`(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)#commit`	Commit candidate configuration to the running configuration
`(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.
`(config-router)#commit`	Commit candidate configuration to the running configuration

© 2026 IP Infusion Inc. Proprietary

R2 Configuration

`(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)#commit`	Commit candidate configuration to the running configuration
`(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.
`(config-router)#commit`	Commit candidate configuration to the running configuration

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 113. Redistribute Routes Into IS-IS

R1 Configuration

`(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)#commit`	Commit candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode.
`(config)#router isis ABC`	Create an IS-IS routing instance for area 49 (ABC).

© 2026 IP Infusion Inc. Proprietary

`(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.
`(config-router)#commit`	Commit candidate configuration to the running configuration

R2 Configuration

`(config)#interface eth1`	Enter interface mode.
`(config-if)#ip address 21.21.21.1/24`	Configure IP address on interface.
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-if)#ip router isis ABC`	Enable IS-IS routing on an interface for area 49 (ABC).
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Enter interface mode.
`(config-if)#ip address 31.31.31.1/24`	Configure IP address on interface.
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration

© 2026 IP Infusion Inc. Proprietary

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

R3 Configuration

`(config)#interface eth2`	Enter interface mode.
`(config-if)#ip address 31.31.31.2/24`	Configure IP address on interface.
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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
`(config-if)#commit`	Commit candidate configuration to the running configuration

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

Users 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`.

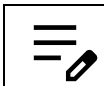

Note: Below configuration is applicable for narrow (non-wide) metric-style. Wide metric can be configured by using the CLI's "`metric-style wide`" under isis instance and "`isis wide-metric < 1-16777214>`" under interface mode.

Topology

Figure 114. Metric Topology

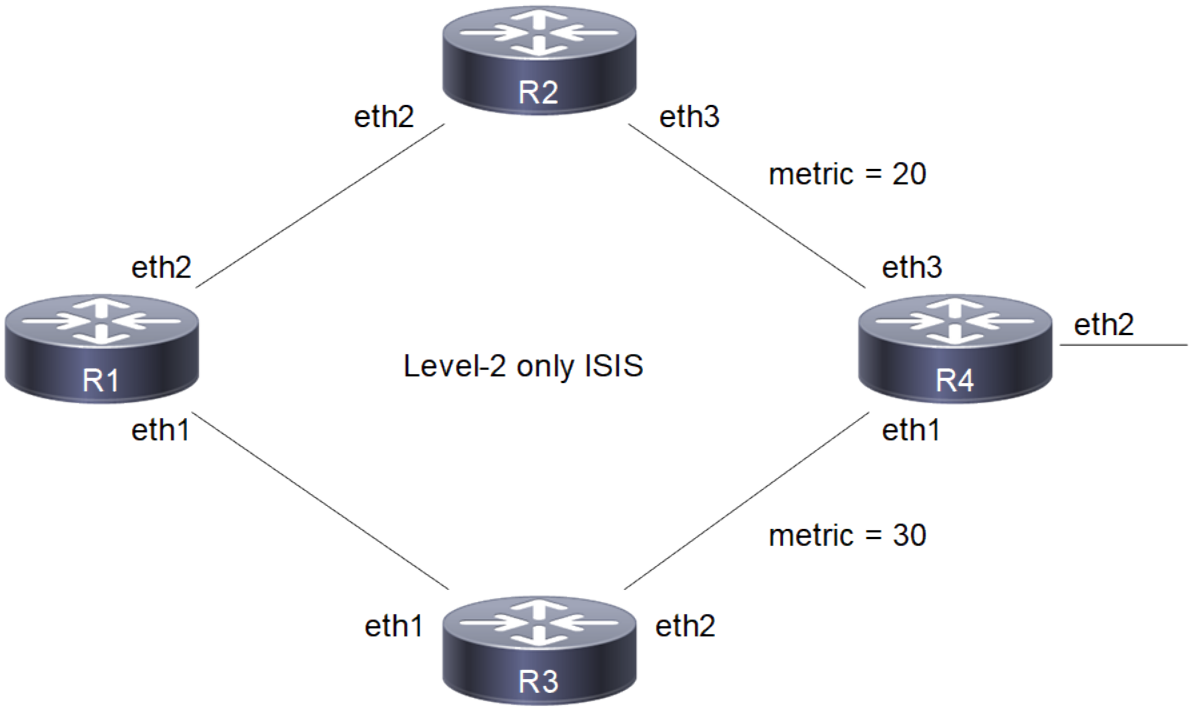

R1 Configuration

`#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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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

© 2026 IP Infusion Inc. Proprietary

	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.
(config-router)#commit	Commit candidate configuration to the running configuration

R2 Configuration

(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)#commit	Commit candidate configuration to the running configuration
(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)#commit	Commit candidate configuration to the running configuration
(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.
(config-router)#commit	Commit candidate configuration to the running configuration

R3 Configuration

(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)#commit	Commit candidate configuration to the running configuration

© 2026 IP Infusion Inc. Proprietary

`(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)#commit`	Commit candidate configuration to the running configuration
`(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 49.0000.0000.0003.00`	Set a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#commit`	Commit candidate configuration to the running configuration

R4 Configuration

`(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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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.

© 2026 IP Infusion Inc. Proprietary

`(config-if)#commit`	Commit candidate configuration to the running configuration

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
```

© 2026 IP Infusion Inc. Proprietary

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 115. Single-Instance L1-L2 Area Routing

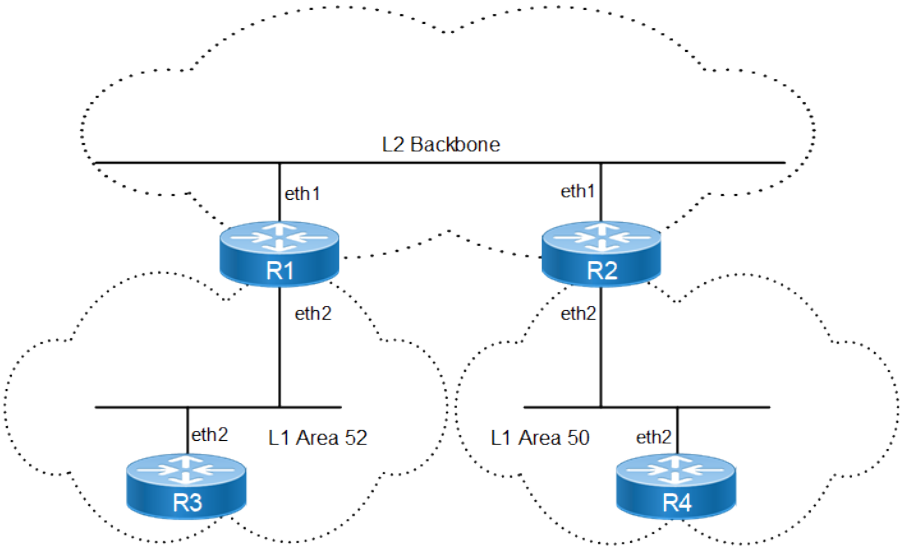

R1 Configuration

`#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.
`(config-if)#isis circuit-type level-2-only`	Set the circuit type for the interface eth1.
`(config-if)#exit`	Exit interface mode.
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(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.

© 2026 IP Infusion Inc. Proprietary

(config-if)#commit	Commit candidate configuration to the running configuration
(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.
(config-router)#commit	Commit candidate configuration to the running configuration

R2 Configuration

(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)#commit	Commit candidate configuration to the running configuration
(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)#commit	Commit candidate configuration to the running configuration
(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.
(config-if)#commit	Commit candidate configuration to the running configuration

R3 Configuration

(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)#commit	Commit candidate configuration to the running configuration

© 2026 IP Infusion Inc. Proprietary

`(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.
`(config-if)#commit`	Commit candidate configuration to the running configuration

R4 Configuration

`(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)#commit`	Commit candidate configuration to the running configuration
`(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.
`(config-if)#commit`	Commit candidate configuration to the running configuration

Validation

```
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
```

© 2026 IP Infusion Inc. Proprietary

```
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 116. Multiple-Instance L1-L2 Area Routing

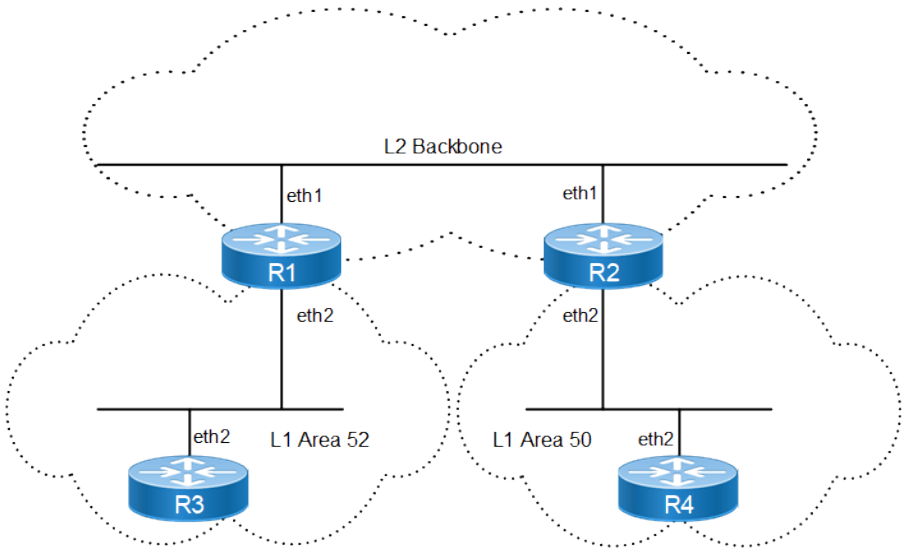

R1 Configuration

`#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)#commit	Commit candidate configuration to the running configuration
(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-if)#commit	Commit candidate configuration to the running configuration
(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)#commit	Commit candidate configuration to the running configuration
(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.
(config-if)#commit	Commit candidate configuration to the running configuration

R2 Configuration

(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)#commit	Commit candidate configuration to the running configuration
(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-if)#commit	Commit candidate configuration to the running

	configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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.
`(config-if)#commit`	Commit candidate configuration to the running configuration

R3 Configuration

`(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 Configuration

`(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.

© 2026 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 117. Route Leaking Topology

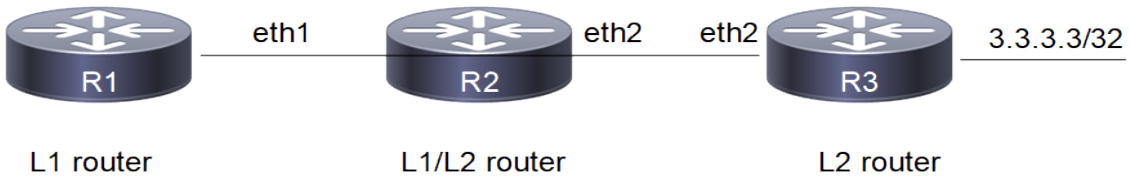

R1 Configuration

`#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-if)#commit`	Commit candidate configuration to the running configuration
`(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).
`(config-if)#commit`	Commit candidate configuration to the running configuration

R2 Configuration

`#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)#commit`	Commit candidate configuration to the running configuration
`(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.

© 2026 IP Infusion Inc. Proprietary

`(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)#commit`	Commit candidate configuration to the running configuration
`(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
`(config-if)#commit`	Commit candidate configuration to the running configuration

R3 Configuration

`#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)#commit`	Commit candidate configuration to the running configuration
`(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-if)#commit`	Commit candidate configuration to the running configuration
`(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).
`(config-if)#commit`	Commit candidate configuration to the running configuration

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.

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

Topology

Figure 118. Route Summarization Topology

R1 Configuration

`#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-if)#commit`	Commit candidate configuration to the running configuration
`(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).
`(config-if)#commit`	Commit candidate configuration to the running configuration

R2 Configuration

`#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)#commit`	Commit candidate configuration to the running configuration
`(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.

© 2026 IP Infusion Inc. Proprietary

(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)#commit	Commit candidate configuration to the running configuration
(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
(config-router-af)#summary-address 66.66.66.0/24 level-1 metric 50	Configure the summary address to summarize IP reachability information.
(config-if)#commit	Commit candidate configuration to the running configuration

R3 Configuration

#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-if)#commit	Commit candidate configuration to the running configuration
(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).
(config-if)#commit	Commit candidate configuration to the running configuration

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 119. IS-IS Distance Topology

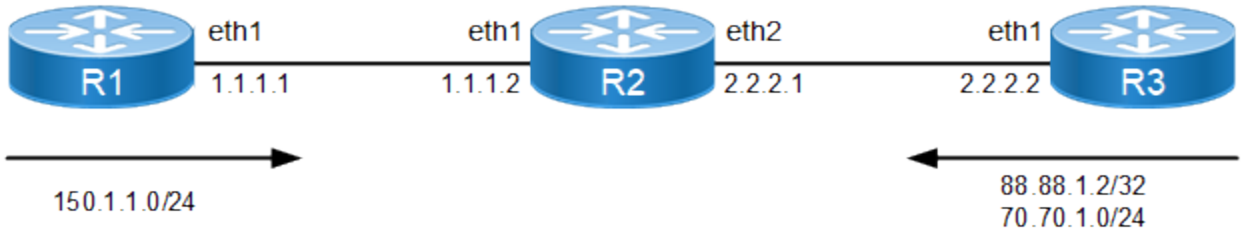

R1 Configuration

`#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)#commit`	Commit candidate configuration to the running configuration
`(config-if)#exit`	
`(config)#ip route 150.1.1.0/24 eth1`	Configure static routes.
`(config)#router isis 1`	Create an IS-IS routing instance (`1`).

© 2026 IP Infusion Inc. Proprietary

(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.
(config-if)#commit	Commit candidate configuration to the running configuration

R2 Configuration

#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-if)#commit	Commit candidate configuration to the running configuration
(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)#commit	Commit candidate configuration to the running configuration
(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)#commit	Commit candidate configuration to the running configuration
(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.
(config-if)#commit	Commit candidate configuration to the running configuration

© 2026 IP Infusion Inc. Proprietary

R3 Configuration

`#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)#commit`	Commit candidate configuration to the running configuration
`(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.
`(config-if)#commit`	Commit candidate configuration to the running configuration

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 [100/10] 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       --
```

© 2026 IP Infusion Inc. Proprietary

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 IS-IS Commands (page 1929) section.

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

The illustration below shows a sample topology with four routers (R1, R2, R3, and R4) connected 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 120. Configure Overload Bit Wait-for-BGP

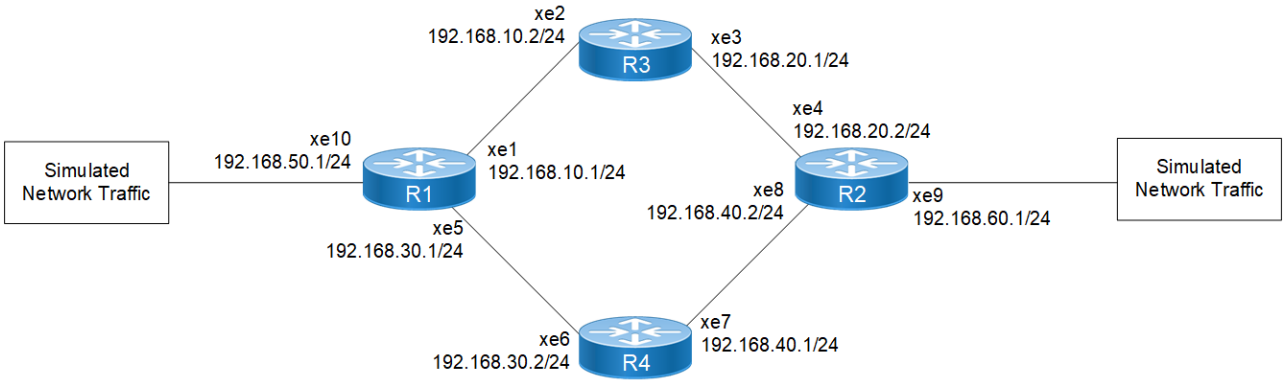

The key element in these configurations is on R3 and R4, where the `set-overload-bit on-startup wait-for-bgp` commands are used. By configuring R3 and R4 with these commands, IS-IS on R3 and R4 will advertise the overload bit to its neighbors, effectively informing them not to route traffic through R3 and R4 until BGP has converged. This behavior helps prevent routing anomalies during network restarts and enhances network stability.

© 2026 IP Infusion Inc. Proprietary

R1 Configuration

The following are the step-by-step configurations on the R1 router.

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.

© 2026 IP Infusion Inc. Proprietary

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 Route-map and Set Weight

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

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.

© 2026 IP Infusion Inc. Proprietary

`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 Configuration

The following are the step-by-step configurations on the R2 router.

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.

© 2026 IP Infusion Inc. Proprietary

`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 Route-map and Set Weight

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

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 Configuration

The following are the step-by-step configurations on the R3 router.

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: Set the Overload Bit

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`.

© 2026 IP Infusion Inc. Proprietary

`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 Configuration

The following are the step-by-step configurations on the R4 router.

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

© 2026 IP Infusion Inc. Proprietary

	running configuration.
`R4(config-if)#exit`	Exits interface mode and returns to the configure mode.

Step 2: Configure IS-IS IPv4 Routing

`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
```

After Reload

```
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.
```

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
```

Remove Neighbor and Reactivate Neighbor

```
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
```

© 2026 IP Infusion Inc. Proprietary

```
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 Inter-level Reachability

The control mechanism of suppressing the redistribution of external, inter-level, or both external and inter-level reachabilities is useful for managing the advertisement of these specific reachability data during overload states, providing flexibility in network configurations.

Topology

The illustration below shows a sample topology with three routers (R1, R2, and R3) interconnected in a linear sequence.

Figure 121. Configure Overload Bit Suppress

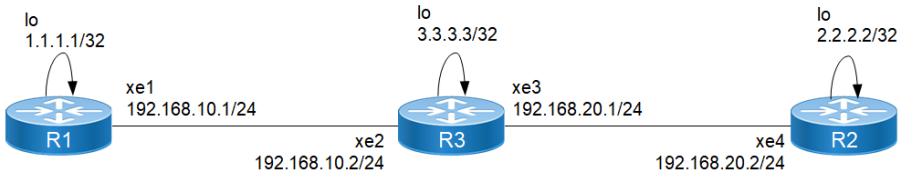

The key element in the below configurations is on R3 router, where the `set-overload-bit suppress external` and `set-overload-bit suppress interlevel` commands are used.

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 IS-IS neighbors or prevents the advertisement of inter-level routes to its IS-IS neighbors. The options `suppress external` or `suppress interlevel` contribute to the controlled management of routing information during network overload states.

R1 Configuration

The following are the step-by-step configurations on the R1 router.

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.

© 2026 IP Infusion Inc. Proprietary

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 inter-level 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`.

© 2026 IP Infusion Inc. Proprietary

`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 Configuration

The following are the step-by-step configurations on the R2 router.

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 Configuration

The following are the step-by-step configurations on the R3 router.

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: Set Overload Bit Options

Configure IS-IS IPv4 routing, setting the overload bit suppress external or inter-level 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 inter-level 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.

© 2026 IP Infusion Inc. Proprietary

Step 4: Configure and Activate BGP Neighbors

> **Note:** Configure BGP only for Suppress External reachability using the `set-overload-bit suppress external` command, excluding inter-level 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: IS-IS Session Status

Before configuring `set-overload-bit suppress external` command, verify the IS-IS 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 IS-IS 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: IS-IS Session Status

After configuring `set-overload-bit suppress external` command, verify the IS-IS 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 IS-IS 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:
    a. The interface loses its adjacency on that interface, for example, `eth1`.
    b. 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:
    a. The interface is IS-IS disabled and must be enabled using the `ip router isis` command (for example, `ip router isis 1`).
    b. If IS-IS is not configured, the interface (for example, eth1) will not be advertised by other IS-IS speaking interfaces to their neighbors.

© 2026 IP Infusion Inc. Proprietary

3. When an interface is configured with the `passive interface` command:
    a. All IS-IS enabled interfaces lose their adjacency.
    b. All IS-IS enabled interfaces in the system will be made passive.
    c. To establish adjacency on a particular interface, the `passive interface <interface-name> disable` command must be enabled.
    d. 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:
    a. All interfaces which are currently passive, will become active.
    b. If IS-IS is configured on those interface, it will start sending out IS-IS packets and attempt to form adjacency.
    c. 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

The illustration below shows a passive-interface configuration example.

Figure 122. IS-IS Passive Interface

R1 Configuration

`#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-if)#commit`	Commit candidate configuration to the running configuration
`(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).
`(config-if)#commit`	Commit candidate configuration to the running configuration

© 2026 IP Infusion Inc. Proprietary

R2 Configuration

`#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)#commit`	Commit candidate configuration to the running configuration
`(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 Configuration

`#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-if)#commit`	Commit candidate configuration to the running configuration
`(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).
`(config-if)#commit`	Commit candidate configuration to the running configuration

© 2026 IP Infusion Inc. Proprietary

Validation

The below validation is a display of the `show clns neighbors` command for R1, R2, and R3 routers configuration.

```
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
```

The below validation is a display of the `show ip isis route` command for R1, R2, and R3 routers configuration.

```
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
```

The below validation is a display of the `show isis database verbose` command for R1, R2, and R3 routers configuration.

```
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
```

The below validation is a display of the `show isis topology` command for R1, R2, and R3 routers configuration.

```
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 Intermediate System to Intermediate System (IS-IS) 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 (page 1925) section.

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-loo ping of packets than a net work 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

The illustration below shows the configuration to enable the basic IS-IS LFA feature.

Figure 123. IS-IS LFA-FRR

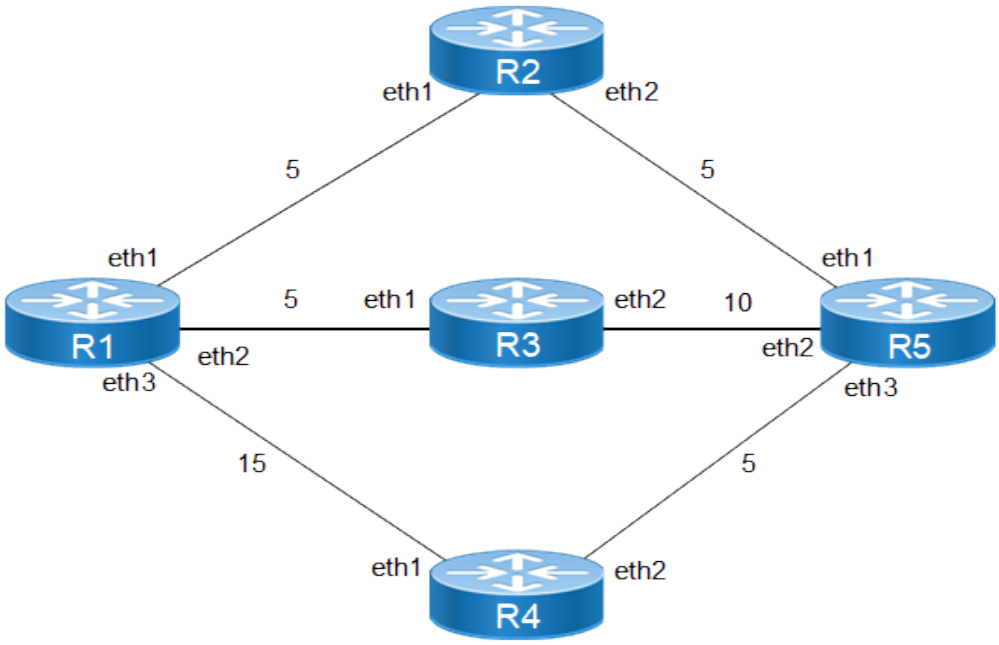

R1 Configuration

`#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 IS-IS 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 IS-IS metric value for interface
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(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 IS-IS 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 IS-IS metric value for interface
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(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 IS-IS 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 IS-IS metric value for interface
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(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 49.0000.000 0.0001.00`	Establish a network entity title, for instance, specifying the area address and the system ID.
`(config-router)#bfd all-interfaces`	Enable BFD for IS-IS on all interfaces
`(config-rou ter)#fast-reroute per-prefix level-1 proto ipv4 all`	Configure LFA-FRR to calculate the available back up path for all L1 IPv4 prefixes learnt
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-rou ter)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

R2 Configuration

`#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 IS-IS 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 IS-IS metric value for interface
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(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 IS-IS 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 IS-IS metric value for interface
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(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 49.0000.000 0.0002.00`	Establish a Network Entity Title for this instance, specifying the area address and the system ID.
`(config-router)#bfd all-interfaces`	Enable BFD for IS-IS on all interfaces
`(config-if)#commit`	Commit candidate configuration to the running configuration

© 2026 IP Infusion Inc. Proprietary

`(config-router)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

R3 Configuration

`#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 IS-IS routing on interface for area 49 with instance 1
`(config-if)#isis metric 5`	Configure IS-IS metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(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 IS-IS routing on interface for area 49 with instance 1
`(config-if)#isis metric 10`	Configure IS-IS metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-if)# exit`	Exit interface mode.
`(config)#router isis 1`	Create an IS-IS routing instance for the area 49 with instance 1
`(config-router)#net 49.0000.000 0.0003.00`	Establish a network entity title, for instance, specifying the area address and the system ID.
`(config-router)#bfd all-interfaces`	Enable BFD for IS-IS on all interfaces
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-rou ter)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

R4 Configuration

`#configure t erminal`	Enter config ure m ode.
`(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 IS-IS routing on interface for area 49 with instance 1
`(config-if)#isis metric 15`	Configure IS-IS metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-if)# exit`	Exit interface mode.

© 2026 IP Infusion Inc. Proprietary

`(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 IS-IS routing on interface for area 49 with instance 1
`(config-if)#isis metric 5`	Configure IS-IS metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-if)# exit`	Exit interface mode.
`(config)#router isis 1`	Create an IS-IS routing instance for the area 49 with instance 1
`(config-router)#net 49.0000.000 0.0004.00`	Establish a network entity title, for instance, specifying the area address and the system ID.
`(config-router)#bfd all-interfaces`	Enable BFD for IS-IS on all interfaces
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-rou ter)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

R5 Configuration

`#configure t erminal`	Enter config ure m ode.
`(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 IS-IS routing on interface for area 49 with instance 1
`(config-if)#isis metric 5`	Configure IS-IS metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(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 IS-IS routing on interface for area 49 with instance 1
`(config-if)#isis metric 10`	Configure IS-IS metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(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 IS-IS routing on interface for area 49 with instance 1
`(config-if)#isis metric 5`	Configure IS-IS metric value for interface
`(config-if)#isis circuit-type level-1`	Enable circuit type on interface

© 2026 IP Infusion Inc. Proprietary

(config-if)#commit	Commit candidate configuration to the running configuration
(config-if)# exit	Exit interface mode.
(config)#router isis 1	Create an IS-IS routing instance for the area 49 with instance 1
(config-router)#net 49.0000.000 0.0005.00	Establish a network entity title, for instance, specifying the area address and the system ID.
(config-router)#bfd all-interfaces	Enable BFD for IS-IS on all interfaces
(config-if)#commit	Commit candidate configuration to the running configuration
(config-rou ter)#exit	Exit router mode.
(config)#exit	Exit config mode.

Validation

IS-IS 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 IS-IS route installation with primary and backup paths in the IS-IS 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
```

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 back up 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
```

Backup Path Based on Route-Map Prefixes

R1 Configuration

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

Apply the Created Route-map with Fast-reroute

`(config)#rou ter isis 1`	Cr eate an IS-IS r out ing 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-if)#commit`	Commit candidate configuration to the running configuration
`(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 Configuration

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 back up 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-if)#commit`	Commit candidate configuration to the running configuration
`(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
```

© 2026 IP Infusion Inc. Proprietary

```
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 Configuration

Configure termination interval on R1 in router mode:

`(config)#rou ter 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-if)#commit`	Commit candidate configuration to the running configuration
`(config-router)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

Validation

R1 Output Before Primary Next-hop Shutdown

```
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
```

Shutdown 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

R1 Output After Primary Next-hop Shutdown

```
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

The illustration below shows the configuration to enable the ISIS LFA feature with ECMP.

Figure 124. ISIS LFA-FRR ECMP

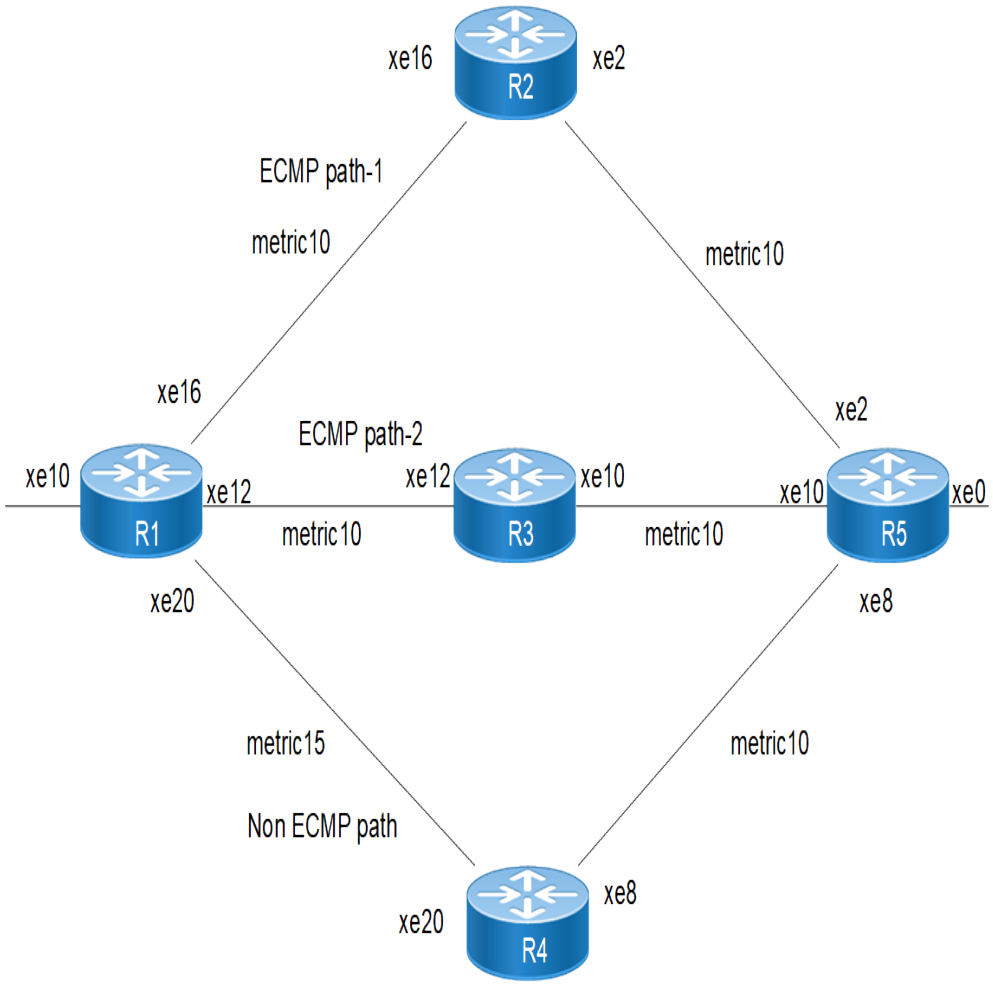

R1 Configuration

`#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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode.
`(config)#router isis 1`	Create an IS-IS routing instance for the area 49 with instance 1
`(config-router)#net 49.0001.0000.0000.0001.00`	Establish a network entity title, for 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`	C onf igure L FA-FRR to c alcul a te the av aila ble back up p ath for all L1 ipv4 prefixes learnt
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-router)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

© 2026 IP Infusion Inc. Proprietary

R2 Configuration

`#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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(config-if)# exit`	Exit interface mode.
`(config)#router isis 1`	Create an IS-IS routing instance for the area 49 with instance 1
`(config-router)#net 49.0001.0000.0000.0002.00`	Establish a network entity title, for instance, specifying the area address and the system ID.
`(config-rout er)#bfd all-interfaces`	Enable BFD for ISIS on all interfaces
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-rou ter)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

R3 Configuration

`#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)#commit`	Commit candidate configuration to the running configuration
`(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.

© 2026 IP Infusion Inc. Proprietary

`(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)#commit`	Commit candidate configuration to the running configuration
`(config-if)# exit`	Exit interface mode.
`(config)#router isis 1`	Create an IS-IS routing instance for the area 49 with instance 1
`(config-router)#net 49.0001.0000.0000.0003.00`	Establish a network entity title, for instance, specifying the area address and the system ID.
`(config-rout er)#bfd all-interfaces`	Enable BFD for ISIS on all interfaces
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-rou ter)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

R4 Configuration

`#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-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)#commit`	Commit candidate configuration to the running configuration
`(config-if)# exit`	Exit interface mode.
`(config)#router isis 1`	Create an IS-IS routing instance for the area 49 with instance 1
`(config-router)#net 49.0001.0000.0000.0004.00`	Establish a network entity title, for instance, specifying the area address and the system ID.
`(config-rout er)#bfd all-interfaces`	Enable BFD for ISIS on all interfaces
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-rou ter)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

R5 Configuration

`#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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode.
`(config)#router isis 1`	Create an IS-IS routing instance for the area 49 with instance 1
`(config-router)#net 49.0001.0000.0000.0005.00`	Establish a network entity title, for instance, specifying the area address and the system ID.
`(config-rout er)#bfd all-interfaces`	Enable BFD for ISIS on all interfaces
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-rou ter)#exit`	Exit router mode.
`(config)#exit`	Exit config mode.

© 2026 IP Infusion Inc. Proprietary

LFA For ECMP Paths Validation

The backup path will be selected by default from same Primary/ECMP set and “Pri” indicates backup selected from ECMP set.

The below validation is a display of the `show clns neighbors` command for R1, R2, R3, R4, and R5 routers configuration.

```
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
```

The validation below displays the fast reroute configuration of the R1 router.

```
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 from primary set itself , backup path will be installed.

`(config)#router isis 1`	Create an IS-IS routing instance for the area 49 with instance 1
`(config-router)#net 49.0000.000 0.0001.00`	Establish a network entity title, for 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)#commit`	Commit candidate configuration to the running configuration
`(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: In the output below, the “Sec” field indicates that 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
```

© 2026 IP Infusion Inc. Proprietary

IS-IS IPv6 Configuration

This section 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 125. Figure 4-46: Basic IS-IS v6 Topology

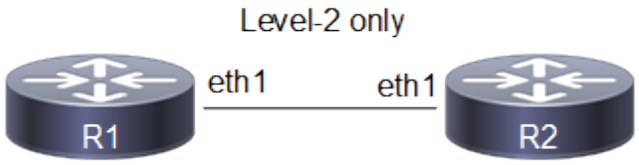

Configuration

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.
`(config-router)#commit`	Commit candidate configuration to the running configuration

R2

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.

© 2026 IP Infusion Inc. Proprietary

`(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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration

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 126. Set IS-IS Priority

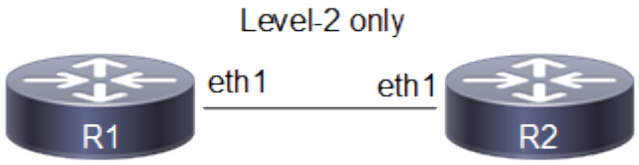

Configuration

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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration

© 2026 IP Infusion Inc. Proprietary

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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration

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 127. Basic dynamic hostname topology

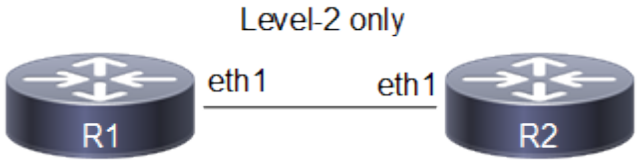

Configuration

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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(config-router)#end`	Exit the current mode and enter privilege mode.

© 2026 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 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 128. 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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration

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)#commit`	Commit candidate configuration to the running configuration

© 2026 IP Infusion Inc. Proprietary

`(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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(config-router)#end`	Exit the current mode and enter privilege mode.

© 2026 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 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.

© 2026 IP Infusion Inc. Proprietary

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 129. Configure IS-IS Metric

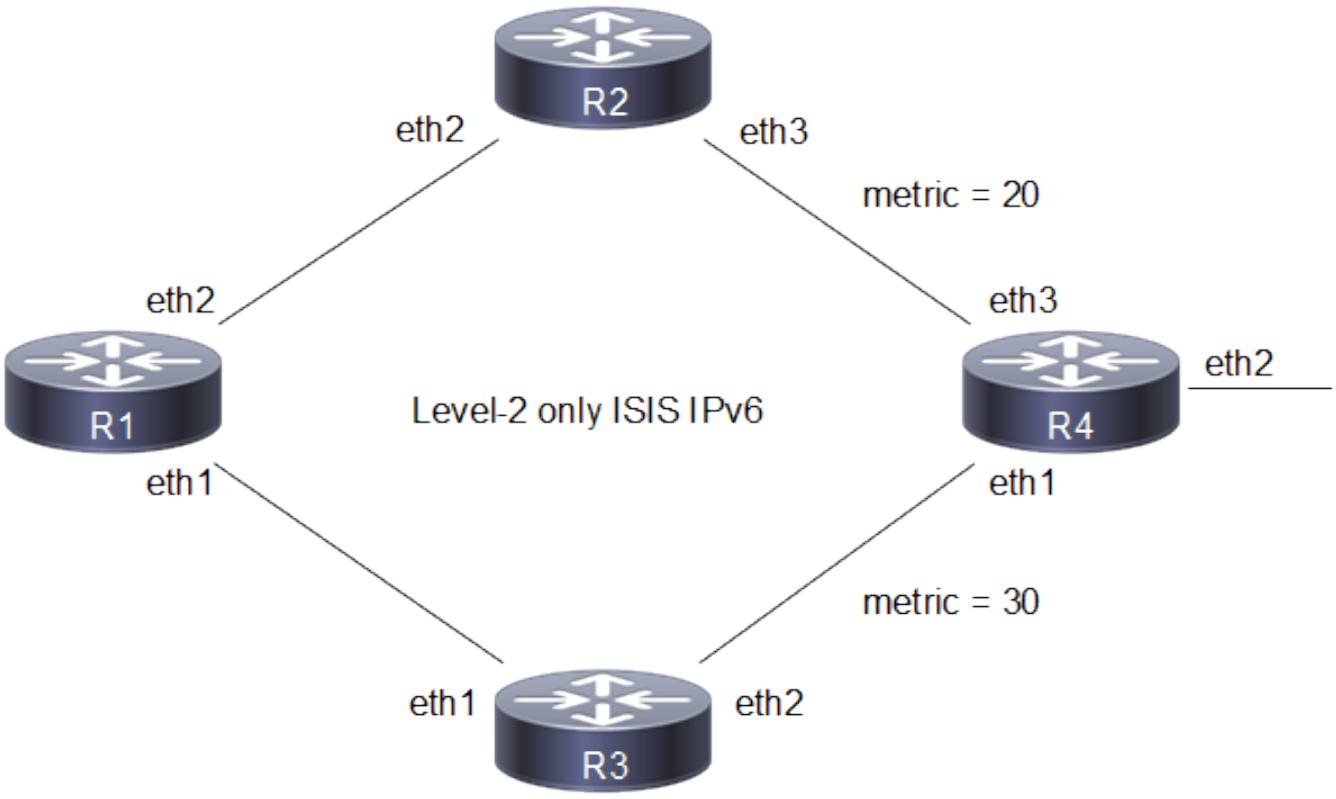

Configuration

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)#commit`	Commit candidate configuration to the running configuration
`(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).

© 2026 IP Infusion Inc. Proprietary

`(config-if)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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.
`(config-router)#commit`	Commit candidate configuration to the running configuration

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

© 2026 IP Infusion Inc. Proprietary

`(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)#commit`	Commit candidate configuration to the running configuration
`(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.
`(config-router)#commit`	Commit candidate configuration to the running configuration

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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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

© 2026 IP Infusion Inc. Proprietary

	process
(config-router-af)#commit	Commit candidate configuration to the running configuration
(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 10:  VRF : default
Total number of routes: 1

C    1111::1/128 [10]
      via ::, lo
```

```
R1#

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 130. Route Summarization 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)#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-router)# is-type level-1`	Configure instance as level-1 routing.
`(config-if)#commit`	Commit candidate configuration to the running

© 2026 IP Infusion Inc. Proprietary

	configuration
(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)#commit	Commit candidate configuration to the running configuration

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)#commit	Commit candidate configuration to the running configuration
(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)#commit	Commit candidate configuration to the running configuration
(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-if)#commit	Commit candidate configuration to the running configuration
(config-router-af)#exit-address-family	Exit address family mode.
(config-router)#exit	Exit router mode.

© 2026 IP Infusion Inc. Proprietary

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 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)#isis circuit-type level-2-only`	Set the circuit type as level-2-only for the interface.
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running configuration
`(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

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

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

R1#show ipv6 route isis
IP Route Table for VRF "default"
i ia    11:1:1:1::/64 [115/68] via fe80::eac5:7aff:fe6b:732e, eth1, 00:00:25
i ia    2000::/64 [115/20] via fe80::eac5:7aff:fe6b:732e, eth1, 00:01:15
R1#

R2#show ipv6 route isis
IP Route Table for VRF "default"
i       11:1:1:1::/64 [115/0] via ::, Null, 00:01:56
i L2    11:1:1:1:1:1::/96 [115/10] via fe80::eac5:7aff:fe7d:5332, eth2, 00:01:56
i L2    11:1:1:1:2:1::/96 [115/10] via fe80::eac5:7aff:fe7d:5332, eth2, 00:01:56
i L2    11:1:1:1:3:1::/96 [115/10] via fe80::eac5:7aff:fe7d:5332, eth2, 00:01:56
R2#

R3#show ipv6 route isis
IP Route Table for VRF "default"
i L2    1000::/64 [115/20] via fe80::eac5:7aff:fe6b:7332, eth2, 00:02:08

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::eac5:7aff:fe6b:732e, eth1
C    1000::/64 [10]
      via ::, xe16
ia   2000::/64 [20]
      via fe80::eac5:7aff:fe6b:732e, eth1
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::eac5:7aff:fe7d:5332, eth2
L2   11:1:1:1:2:1::/96 [10]
      via fe80::eac5:7aff:fe7d:5332, eth2
L2   11:1:1:1:3:1::/96 [10]
      via fe80::eac5:7aff:fe7d:5332, eth2
C    1000::/64 [10]
      via ::, xe16
C    2000::/64 [10]
      via ::, xe20

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 ::
L2   1000::/64 [20]
```

```
      via fe80::eac5:7aff:fe6b:7332, eth2
C    2000::/64 [10]
      via ::, eth2

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   0x4755        686               0/0/0
  Area Address: 49.0001
  NLPID:        0x8E
  IPv6 Address: 1000::1
  Metric:   10         IS 0000.0000.0002.01
  Metric:   10         IPv6 1000::/64
0000.0000.0002.00-00  0x00000004   0x4A36        713               0/0/0
  Area Address: 49.0001
  NLPID:        0x8E
  IPv6 Address: 1000::2
  Metric:   10         IS 0000.0000.0002.01
  Metric:   10         IPv6 1000::/64
  Metric:   10         IPv6-Interarea 2000::/64
  Metric:   58         IPv6-Interarea 11:1:1:1::/64
0000.0000.0002.01-00  0x00000001   0x0ECB        685               0/0/0
  Metric:   0          IS 0000.0000.0002.00
  Metric:   0          IS 0000.0000.0001.00

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   0x4755        657               0/0/0
  Area Address: 49.0001
  NLPID:        0x8E
  IPv6 Address: 1000::1
  Metric:   10         IS 0000.0000.0002.01
  Metric:   10         IPv6 1000::/64
0000.0000.0002.00-00* 0x00000004   0x4A36        686               0/0/0
  Area Address: 49.0001
  NLPID:        0x8E
  IPv6 Address: 1000::2
  Metric:   10         IS 0000.0000.0002.01
  Metric:   10         IPv6 1000::/64
  Metric:   10         IPv6-Interarea 2000::/64
  Metric:   58         IPv6-Interarea 11:1:1:1::/64
0000.0000.0002.01-00* 0x00000001   0x0ECB        658               0/0/0
  Metric:   0          IS 0000.0000.0002.00
  Metric:   0          IS 0000.0000.0001.00

IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
0000.0000.0002.00-00* 0x00000004   0x67A6        676               0/0/0
  Area Address: 49.0001
  NLPID:        0x8E
  IPv6 Address: 2000::1
  Metric:   10         IS 0000.0000.0003.01
  Metric:   10         IPv6 2000::/64
  Metric:   10         IPv6 1000::/64
0000.0000.0003.00-00  0x00000002   0xC118        675               0/0/0
  Area Address: 49.0001
  NLPID:        0x8E
  IPv6 Address: 2000::2
  Metric:   10         IS 0000.0000.0003.01
  Metric:   10         IPv6 2000::/64
  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        671               0/0/0
  Metric:   0          IS 0000.0000.0003.00
```

```
  Metric:   0          IS 0000.0000.0002.00

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
0000.0000.0002.00-00  0x00000004   0x67A6        605               0/0/0
  Area Address: 49.0001
  NLPID:        0x8E
  IPv6 Address: 2000::1
  Metric:   10         IS 0000.0000.0003.01
  Metric:   10         IPv6 2000::/64
  Metric:   10         IPv6 1000::/64
0000.0000.0003.00-00* 0x00000002   0xC118        606               0/0/0
  Area Address: 49.0001
  NLPID:        0x8E
  IPv6 Address: 2000::2
  Metric:   10         IS 0000.0000.0003.01
  Metric:   10         IPv6 2000::/64
  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        602               0/0/0
  Metric:   0          IS 0000.0000.0003.00
  Metric:   0          IS 0000.0000.0002.00
```

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.)

© 2026 IP Infusion Inc. Proprietary

- 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 131. 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)#commit`	Commit candidate configuration to the running configuration
`(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-if)#commit`	Commit candidate configuration to the running configuration
`(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)#commit`	Commit candidate configuration to the running

	configuration
(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)#commit	Commit candidate configuration to the running configuration
(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-if)#commit	Commit candidate configuration to the running configuration
(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)#commit	Commit candidate configuration to the running configuration
(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-if)#commit	Commit candidate configuration to the running configuration
(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 132. Basic BFD over IS-ISv6 Topology

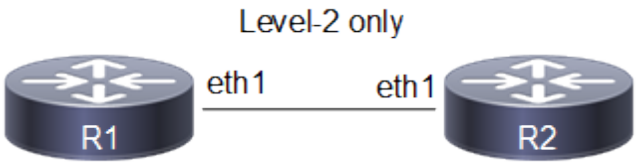

Configuration

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)#commit`	Commit candidate configuration to the running configuration
`(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

© 2026 IP Infusion Inc. Proprietary

	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)#commit	Commit candidate configuration to the running configuration
(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)#commit	Commit candidate configuration to the running configuration
(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)#commit	Commit candidate configuration to the running configuration
(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 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 133. Basic IS-ISv6 Topology

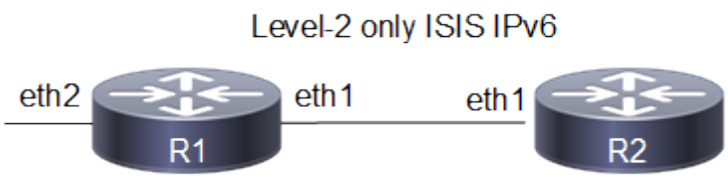

Configuration

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)#commit	Commit candidate configuration to the running configuration
(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)#commit	Commit candidate configuration to the running configuration
(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)#commit	Commit candidate configuration to the running configuration
(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)#commit	Commit candidate configuration to the running configuration
(config-router)#end	Exit current mode and enter privilege mode

© 2026 IP Infusion Inc. Proprietary

Validation

```
R1#sh 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         cc37.ab87.3a74      Up     23        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     6         L2   IS-IS
R2#

R1#show ipv6 route isis
IP Route Table for VRF "default"
R1#

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
E    ::/0 [0]
      via ::
R1#

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
```

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 IPv6 address family.

Topology

Figure 134. IS-IS Distance Topology

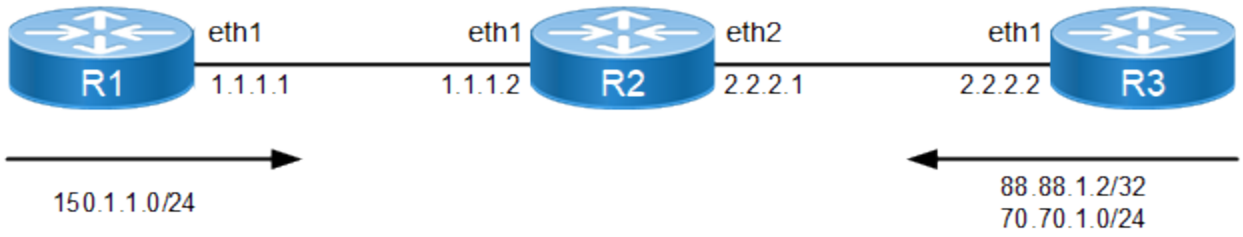

Configuration

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)#ipv6 address 1000::1/64`	Configure IPv6 address on interface eth1.
`(config)#ip router isis 1`	Create an IS-IS routing instance on interface eth1.
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-if)#exit`	
`(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.
`(config-if)#commit`	Commit candidate configuration to the running configuration

R2

`#configure terminal`	Enter configure mode
`(config)#ipv6 access-list DIST`	Enter access list mode
`(config-ip-acl)#permit ipipv6 6000::/64 any any`	Create an access list to permit the 88.88.1.2/32 route from R3.
`(config-if)#commit`	Commit candidate configuration to the running

© 2026 IP Infusion Inc. Proprietary

	configuration
(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)#ipv6 address 1000::2/64	Configure IPv6 address on interface.
(config-if)#ip router isis 1	Enable IS-IS routing on interface eth1
(config-if)#commit	Commit candidate configuration to the running configuration
(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)#ipv6 address 4000::1/64	Configure IPv6 address on interface.
(config-if)#ip router isis 1	Enable IS-IS routing on interface eth2
(config-if)#commit	Commit candidate configuration to the running configuration
(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)#address-family ipv6	Enter 'address-family ipv6' mode, where users can configure IPv6 routing specific configuration
(config-router-af)#distance 100	Configure the administrative distance for all routes received from R1 and R2.
(config-router-af)#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-af)#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.
(config-if)#commit	Commit candidate configuration to the running configuration

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)#commit	Commit candidate configuration to the running

	configuration
`(config-if)#exit`	Exit interface mode.
`(config)#ipv6 route 5000::/64 eth2`	Configure static routes
`(config)#ipv6 route 6000::/64 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.
`(config-router)#address-family ipv6`	Enter 'address-family ipv6' mode, where users can configure IPv6 routing specific configuration
`(config-router-af)#redistribute static`	Redistribute the static routes
`(config-if)#commit`	Commit candidate configuration to the running configuration

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

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
0001.0000.0003 eth2 5254.00a8.940d Up 8 L1 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
```

© 2026 IP Infusion Inc. Proprietary

IS-IS IPv4 TE-Metric Extension

Overview

This section contains basic IS-IS TE-Metric Extension configuration examples.

A TE link represents an IS-IS/OSPF link state advertisement and a link state database of certain physical resources and their properties between two nodes. Typically, a TE link is advertised as an adjunct to a "regular" OSPF or IS-IS link. That is, an adjacency is brought up on the link. When the link is up, both the regular IGP properties of the link (for example, the SPF metric) and the TE properties of the link are then advertised.

Note: Bandwidth-measurement values should be less than maximum link bandwidth configuration.

Topology

Figure 135. Basic IS-IS Topology

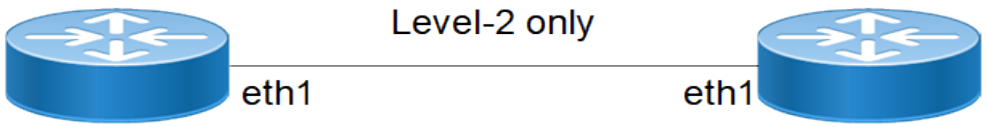

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode.
`(config-if)#ip address 21.21.21.2/24`	Configure IP address on interface.
`(config-if)#isis circuit-type level-2`	Configure ISIS circuit-type as Level-2.
`(config-if)#isis te-metric 40 level-2`	Configure ISIS te-metric
`(config-if)#bandwidth-measurement static uni-available-bandwidth 8g`	Configure Bandwidth-measurement uni-available-bandwidth
`(config-if)#delay-measurement static uni-delay-variation 12`	Configure delay-measurement uni-delay-variation static
`(config-if)#delay-measurement static uni-link-delay 12`	Configure delay-measurement uni-link-delay static
`(config-if)#loss-measurement static uni-link-loss 12.3`	Configure loss-measurement uni-link-loss static
`(config-if)#ip router isis ABC`	Enable IS-IS routing on an interface for area 49 (ABC).
`(config-if)#commit`	Commit candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

(config-if)#exit	Exit interface mode.
(config)#router isis ABC	Create an IS-IS routing instance for area 49 (ABC).
(config-router)#metric-style wide	Configure the new style of metric type as wide.
(config-router)#mpls traffic-eng level-2	Enable MPLS-TE in is-type Level-2
(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)#commit	Commit candidate configuration to the running configuration.

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)#isis circuit-type level-2	Configure ISIS circuit-type as Level-2.
(config-if)#isis te-metric 40 level-2	Configure ISIS Te-metric
(config-if)#delay-measurement static uni-delay-variation 12	Configure delay-measurement uni-delay-variation static
(config-if)#delay-measurement static uni-link-delay 24	Configure delay-measurement uni-link-delay static
(config-if)#loss-measurement static uni-link-loss 23.3	Configure loss-measurement uni-link-loss static
(config-if)#commit	Commit candidate configuration to the running configuration.
(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)#metric-style wide	Configure the new style of metric type as wide.
(config-router)#mpls traffic-eng level-2	Enable MPLS-TE in is-type Level-2
(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)#commit	Commit candidate configuration to the running configuration.

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 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* 0x00000091   0x6A2D        1153              0/0/0
  Area Address: 49
  NLPID:        0xCC
  IP Address:   21.21.21.2
  Router ID:    10.12.65.109
  Metric:   10         IS-Extended 0000.0000.0001.01
    IPv4 Interface Address: 21.21.21.2
    Neighbor IP Address: 21.21.21.2
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 40
    Link Delay : 12 us  Anomalous : 0
    Link Delay-variation : 12 us
    Link Loss : 12.300000%, Anomalous : 0
    Link Available Bandwidth: 8g
  Metric:   10         IP-Extended 21.21.21.0/24
0000.0000.0001.01-00* 0x0000008D   0x78C5        628               0/0/0
  Metric:   0          IS-Extended 0000.0000.0001.00
  Metric:   0          IS-Extended 0000.0000.0002.00
0000.0000.0002.00-00  0x0000008F   0xABF0        992               0/0/0
  Area Address: 49
  NLPID:        0xCC
  IP Address:   21.21.21.1
  Router ID:    10.12.65.153
  Metric:   10         IS-Extended 0000.0000.0001.01
    IPv4 Interface Address: 21.21.21.1
    Neighbor IP Address: 21.21.21.2
    Maximum Link Bandwidth: 100m
```

```
    Reservable Bandwidth: 100m
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100m
      Unreserved Bandwidth at priority 1: 100m
      Unreserved Bandwidth at priority 2: 100m
      Unreserved Bandwidth at priority 3: 100m
      Unreserved Bandwidth at priority 4: 100m
      Unreserved Bandwidth at priority 5: 100m
      Unreserved Bandwidth at priority 6: 100m
      Unreserved Bandwidth at priority 7: 100m
    TE-Default Metric: 30
    Link Delay : 24 us  Anomalous : 0
    Link Delay-variation : 12 us
    Link Loss : 23.299998%, Anomalous : 0
  Metric:   10         IP-Extended 21.21.21.0/24
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  0x00000092   0x682E        833               0/0/0
  Area Address: 49
  NLPID:        0xCC
  IP Address:   21.21.21.2
  Router ID:    10.12.65.109
  Metric:   10         IS-Extended 0000.0000.0001.01
    IPv4 Interface Address: 21.21.21.2
    Neighbor IP Address: 21.21.21.2
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 40
    Link Delay : 12 us  Anomalous : 0
    Link Delay-variation : 12 us
    Link Loss : 12.300000%, Anomalous : 0
    Link Available Bandwidth: 8g
  Metric:   10         IP-Extended 21.21.21.0/24
0000.0000.0001.01-00  0x0000008E   0x76C6        308               0/0/0
  Metric:   0          IS-Extended 0000.0000.0001.00
  Metric:   0          IS-Extended 0000.0000.0002.00
0000.0000.0002.00-00* 0x00000090   0xA9F1        674               0/0/0
  Area Address: 49
  NLPID:        0xCC
  IP Address:   21.21.21.1
  Router ID:    10.12.65.153
  Metric:   10         IS-Extended 0000.0000.0001.01
    IPv4 Interface Address: 21.21.21.1
    Neighbor IP Address: 21.21.21.2
    Maximum Link Bandwidth: 100m
    Reservable Bandwidth: 100m
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100m
      Unreserved Bandwidth at priority 1: 100m
      Unreserved Bandwidth at priority 2: 100m
      Unreserved Bandwidth at priority 3: 100m
      Unreserved Bandwidth at priority 4: 100m
      Unreserved Bandwidth at priority 5: 100m
      Unreserved Bandwidth at priority 6: 100m
      Unreserved Bandwidth at priority 7: 100m
    TE-Default Metric: 30
    Link Delay : 24 us  Anomalous : 0
```

```
    Link Delay-variation : 12 us
    Link Loss : 23.299998%, Anomalous : 0
  Metric:   10         IP-Extended 21.21.21.0/24
```

IS-IS Graceful Restart Configuration

Overview

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 is capable to restart gracefully with non-stop forwarding during the recovery. And the Helper ISIS router is 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 136. IS-IS Graceful Restart

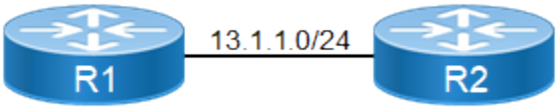

Configuration

The following configuration assumes 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

© 2026 IP Infusion Inc. Proprietary

R2

`#configure terminal`	Enter configure mode.
`(config)#isis restart helper`	Configure this router as a restart helper.
`(config)#commit`	Commit candidate configuration to the running configuration

Notes:

- The IS-IS daemon in the restarting router must be manually restarted using the restart IS-IS graceful command; it does not restart automatically.
- Execute `restart isis graceful grace-period <1-65535>` CLI in privilege mode directly to view the graceful restart warning message.
- The IS-IS daemon in the restarting router must be manually restarted using the restart IS-IS graceful command; it does not restart automatically.
- The unplanned GR is not supported.
- 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
```

© 2026 IP Infusion Inc. Proprietary

```
       * - 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
```

IS-IS-TE IPv4

This section contains basic IS-IS TE configuration examples.

A TE link represents an IS-IS/OSPF link state advertisement and a link state database of certain physical resources and their properties between two nodes. Typically, a TE link is advertised as an adjunct to a “regular” OSPF or IS-IS link. That is, an adjacency is brought up on the link. When the link is up, both the regular IGP properties of the link (for example, the SPF metric) and the TE properties of the link are then advertised.

Enable MPLS-TE in Level-1 and Level-2 on L1-L2 IS

In the following example, R1 is the L1 router, R2 is the L1/L2 router enabling MPLS-TE for both Level-1 and Level-2 IS, and R3 is the L2 router. The following configuration is given for R1, R2 and R3.

Toplogy

Figure 137. MPLS-TE Topology

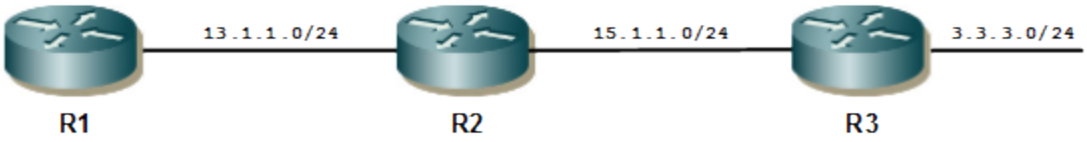

Configuration

R2

`#configure terminal`	Enter Configure mode.
`(config)#interface eth1`	Specify the interface (`eth1`) to configure, and enter Interface mode.
`(config-if)#ip address 13.1.1.2/24`	Configure ip address to the interface eth1.
`(config-if)#ip router isis 1`	Enable ISIS on interface eth1.
`(config-if)#isis circuit-type level-1`	Configure ISIS circuit-type as Level-1
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-if)#exit`	Exit from interface configuration mode.
`(config)#interface eth2`	Enter Interface eth2 configure mode.
`(config-if)#ip address 15.1.1.2/24`	Configure ip address to the interface eth2.
`(config-if)#ip router isis 1`	Enable ISIS on interface eth2.
`(config-if)#isis circuit-type level-2`	Configure ISIS circuit-type as Level-2.
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-if)#exit`	Exit from interface configuration mode.
`(config)#router isis 1`	Create an IS-IS routing instance for area 49 (1).

© 2026 IP Infusion Inc. Proprietary

(config-router)#metric-style wide	Configure the new style of metric type as wide.
(config-router)#mpls traffic-eng router-id 2.2.2.2	Configure MPLS-TE unique router-id TLV.
(config-router)#mpls traffic-eng level-1	Enable MPLS-TE in is-type Level-1.
(config-router)#mpls traffic-eng level-2	Enable MPLS-TE in is-type Level-2.
(config-router)#net 49.0001.0000.0000.0002.00	Configure the ISIS net address with area id as: 49.0001 and system id as: 0000.0000.0002.
(config-router)#commit	Commit candidate configuration to the running configuration
(config-router)#end	Exit Router mode.

R1

#configure terminal	Enter Configure mode.
(config)#interface eth1	Specify the interface (eth1) to configure, and enter Interface mode.
(config-if)#ip address 13.1.1.1/24	Configure ip address to the interface eth1.
(config-if)#ip router isis 1	Enable ISIS on interface eth1.
(config-if)#isis circuit-type level-1	Configure ISIS circuit-type as Level-1
(config-if)#commit	Commit candidate configuration to the running configuration
(config-if)#exit	Exit from interface configuration mode.
(config)#router isis 1	Create an IS-IS routing instance for area 49 (1).
(config-router)#metric-style wide	Configure the new style of metric type as wide.
(config-router)#is-type level-1	Configure IS-Type as Level-1 router.
(config-router)#mpls traffic-eng router-id 1.1.1.1	Configure MPLS-TE unique router-id TLV.
(config-router)#mpls traffic-eng level-1	Enable MPLS-TE in is-type Level-1.
(config-router)#net 49.0001.0000.0000.0001.00	Configure the ISIS net address with area id as: 49.0001 and system-id as: 0000.0000.0001.
(config-router)#commit	Commit candidate configuration to the running configuration
(config-router)#end	Exit Router mode.

R3

#configure terminal	Enter Configure mode.
(config)#interface eth2	Specify the interface (eth2) to configure, and enter Interface mode.
(config-if)#ip address 15.1.1.3/24	Configure ip address to the interface eth1.
(config-if)#ip router isis 1	Enable ISIS on interface eth1.
(config-if)#isis circuit-type level-2	Configure ISIS circuit-type as Level-2

© 2026 IP Infusion Inc. Proprietary

(config-if)#commit	Commit candidate configuration to the running configuration
(config-if)#exit	Exit from interface configuration mode.
(config)#router isis 1	Create an IS-IS routing instance for area 49 (1).
(config-router)#metric-style wide	Configure the new style of metric type as wide.
(config-router)#is-type level-2	Configure IS-Type as Level-1 router.
(config-router)#mpls traffic-eng router-id 3.3.3.3	Configure MPLS-TE unique router-id TLV.
(config-router)#mpls traffic-eng level-2	Enable MPLS-TE in is-type Level-2.
(config-router)#net 49.0001.0000.0000.0003.00	Configure the ISIS net address with area id as: 49.0001 and system-id as: 0000.0000.0003.
(config-router)#commit	Commit candidate configuration to the running configuration
(config-router)#end	Exit Router mode.

Validation

R2

Check the output of show clns neighbors (page 2034) to verify that ISIS adjacency is up.

```
#show clns neighbors
Area 1:
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        5254.00bb.5e85      Up     24        L1   IS-IS
0000.0000.0003 eth2        5254.00ac.f960      Up     9         L2   IS-IS
#
```

Check the output of **show isis database (page 2056)level-1 verbose** to verify that LSP does have correct router-id.

```
R2#show isis database level-1 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  0x00000003   0x5D0E        629               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   13.1.1.1
  Router ID:    1.1.1.1
  Metric:   10         IS-Extended 0000.0000.0001.01
    IPv4 Interface Address: 13.1.1.1
    Neighbor IP Address: 13.1.1.1
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 10
  Metric:   10         IP-Extended 13.1.1.0/24
0000.0000.0001.01-00  0x00000002   0x8D3E        629               0/0/0
```

```
  Metric:   0          IS-Extended 0000.0000.0001.00
  Metric:   0          IS-Extended 0000.0000.0002.00
0000.0000.0002.00-00* 0x00000004   0xE080        634               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   13.1.1.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended 0000.0000.0001.01
    IPv4 Interface Address: 13.1.1.2
    Neighbor IP Address: 13.1.1.1
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 10
  Metric:   10         IP-Extended 13.1.1.0/24
```

Check the output of **show isis database (page 2056)level-2 verbose** to verify that LSP does have correct router-id.

```
R2#show isis database level-2 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   0xEE2C        622               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   15.1.1.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended 0000.0000.0003.01
    IPv4 Interface Address: 15.1.1.2
    Neighbor IP Address: 15.1.1.3
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 10
  Metric:   10         IP-Extended 15.1.1.0/24
  Metric:   10         IP-Extended 13.1.1.0/24
0000.0000.0003.00-00  0x00000003   0x94BA        616               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   15.1.1.3
  Router ID:    3.3.3.3
  Metric:   10         IS-Extended 0000.0000.0003.01
    IPv4 Interface Address: 15.1.1.3
    Neighbor IP Address: 15.1.1.3
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
```

```
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 10
  Metric:   10         IP-Extended 15.1.1.0/24
0000.0000.0003.01-00  0x00000002   0x9134        616               0/0/0
  Metric:   0          IS-Extended 0000.0000.0003.00
  Metric:   0          IS-Extended 0000.0000.0002.00
```

R1

Check the output of show clns neighbors (page 2034) to verify that ISIS adjacency is up.

```
#show clns neighbors

Area 1:
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        5254.00f6.4ae7  Up     7         L1   IS-IS
```

Check the output of **show isis database (page 2056)level-1 verbose** to verify that LSP does have correct router-id.

```
R1#show isis database level-1 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* 0x00000003   0x5D0E        364               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   13.1.1.1
  Router ID:    1.1.1.1
  Metric:   10         IS-Extended 0000.0000.0001.01
    IPv4 Interface Address: 13.1.1.1
    Neighbor IP Address: 13.1.1.1
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 10
  Metric:   10         IP-Extended 13.1.1.0/24
0000.0000.0001.01-00* 0x00000002   0x8D3E        364               0/0/0
  Metric:   0          IS-Extended 0000.0000.0001.00
  Metric:   0          IS-Extended 0000.0000.0002.00
0000.0000.0002.00-00  0x00000004   0xE080        368               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   13.1.1.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended 0000.0000.0001.01
    IPv4 Interface Address: 13.1.1.2
    Neighbor IP Address: 13.1.1.1
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
```

```
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 10
  Metric:   10         IP-Extended 13.1.1.0/24
```

R3

Check the output of show clns neighbors (page 2034) to verify that ISIS adjacency is up.

```
#sh clns  neighbors

Area 1:
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth2        5254.0062.3ea6      Up     21        L2   IS-IS
```

Check the output of **show isis database (page 2056)level-2 verbose** to verify that LSP does have correct router-id.

```
R3#show isis database level-2 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   0xEE2C        337               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   15.1.1.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended 0000.0000.0003.01
    IPv4 Interface Address: 15.1.1.2
    Neighbor IP Address: 15.1.1.3
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 10
  Metric:   10         IP-Extended 15.1.1.0/24
  Metric:   10         IP-Extended 13.1.1.0/24
0000.0000.0003.00-00* 0x00000003   0x94BA        333               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   15.1.1.3
  Router ID:    3.3.3.3
  Metric:   10         IS-Extended 0000.0000.0003.01
    IPv4 Interface Address: 15.1.1.3
    Neighbor IP Address: 15.1.1.3
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
```

```
    TE-Default Metric: 10
  Metric:   10          IP-Extended 15.1.1.0/24
0000.0000.0003.01-00* 0x00000002   0x9134        333               0/0/0
  Metric:   0           IS-Extended 0000.0000.0003.00
  Metric:   0           IS-Extended 0000.0000.0002.00
```

Maximum Link Bandwidth and Reservable Bandwidth

In the following example, R1 is the L1 router, R2 is the L1/L2 router enabling MPLS-TE for both Level-1 and Level-2 IS, and R3 is the L2 router. The following configuration is given for R1, R2 and R3.

Topology

Figure 138. MPLS-TE Topology

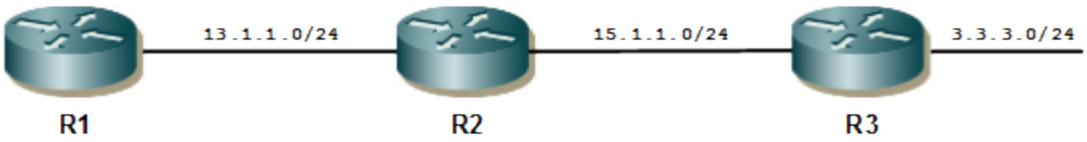

Configuration

R2

`#configure terminal`	Enter Configure mode.
`(config)#interface eth1`	Specify the interface (`eth1`) to configure, and enter Interface mode.
`(config-if)#ip address 13.1.1.2/24`	Configure ip address to the interface eth1.
`(config-if)#ip router isis 1`	Enable ISIS on interface eth1.
`(config-if)#isis circuit-type level-1`	Configure ISIS circuit-type as Level-1
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-if)#exit`	Exit from interface configuration mode.
`(config-if)#interface eth2`	Enter Interface eth2 configure mode.
`(config-if)#ip address 15.1.1.2/24`	Configure ip address to the interface eth2.
`(config-if)#ip router isis 1`	Enable ISIS on interface eth2.
`(config-if)#isis circuit-type level-2`	Configure ISIS circuit-type as Level-2.
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-if)#exit`	Exit from interface configuration mode.
`(config)#router isis 1`	Create an IS-IS routing instance for area 49 (1).
`(config-router)#metric-style wide`	Configure the new style of metric type as wide.
`(config-router)#mpls traffic-eng router-id 2.2.2.2`	Configure MPLS-TE unique router-id TLV.
`(config-router)#mpls traffic-eng level-1`	Enable MPLS-TE in is-type Level-1.

© 2026 IP Infusion Inc. Proprietary

`(config-router)#mpls traffic-eng level-2`	Enable MPLS-TE in is-type Level-2.
`(config-router)#net 49.0001.0000.0000.0002.00`	Configure the ISIS net address with area id as: 49.0001 and system id as: 0000.0000.0002.
`(config-router)#commit`	Commit candidate configuration to the running configuration
`(config-router)#exit`	Exit Router mode.
`(config)#interface eth2`	Enter the interface configure mode.
`(config-if)#bandwidth 1g`	Configure Maximum link Bandwidth as 1g bits per second.
`(config-if)#reservable-bandwidth 100m`	Specify the maximum reservable bandwidth per interface.
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-if)#exit`	Exit interface configure mode.
`(config)#interface eth1`	Enter the interface configure mode.
`(config-if)#bandwidth 500m`	Configure Maximum link Bandwidth as 500m bits per second.
`(config-if)#reservable-bandwidth 200m`	Specify the maximum reservable bandwidth per interface.
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-if)#exit`	Exit interface configure mode.

R1

`#configure terminal`	Enter Configure mode.
`(config)#interface eth1`	Specify the interface (`eth1`) to configure, and enter Interface mode.
`(config-if)#ip address 13.1.1.1/24`	Configure ip address to the interface eth1.
`(config-if)#ip router isis 1`	Enable ISIS on interface eth1.
`(config-if)#isis circuit-type level-1`	Configure ISIS circuit-type as Level-1
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-if)#exit`	Exit from interface configuration mode.
`(config)#router isis 1`	Create an IS-IS routing instance for area 49 (1).
`(config-router)#metric-style wide`	Configure the new style of metric type as wide.
`(config-router)#is-type level-1`	Configure IS-Type as Level-1 router.
`(config-router)#mpls traffic-eng router-id 1.1.1.1`	
`(config-router)#mpls traffic-eng level-1`	Enable MPLS-TE in is-type Level-1.
`(config-router)#net 49.0001.0000.0000.0001.00`	Configure the ISIS net address with area id as:

	49.0001 and system-id as: 0000.0000.0001.
(config-router)#commit	Commit candidate configuration to the running configuration
(config-router)#exit	Exit Router mode.
(config)#interface eth1	Enter the interface configure mode.
(config-if)#bandwidth 1g	Configure Maximum link Bandwidth as 1g bits per second.
(config-if)#reservable-bandwidth 100m	Specify the maximum reservable bandwidth per interface.
(config-if)#commit	Commit candidate configuration to the running configuration
(config-if)#exit	Exit interface configure mode.

R3

#configure terminal	Enter Configure mode.
(config)#interface eth2	Specify the interface (eth2) to configure, and enter Interface mode.
(config-if)#ip address 15.1.1.3/24	Configure ip address to the interface eth1.
(config-if)#ip router isis 1	Enable ISIS on interface eth1.
(config-if)#isis circuit-type level-2	Configure ISIS circuit-type as Level-1
(config-if)#commit	Commit candidate configuration to the running configuration
(config-if)#exit	Exit from interface configuration mode.
(config)#router isis 1	Create an IS-IS routing instance for area 49 (1).
(config-router)#metric-style wide	Configure the new style of metric type as wide.
(config-router)#is-type level-2	Configure IS-Type as Level-2 router.
(config-router)#mpls traffic-eng router-id 3.3.3.3	Configure MPLS-TE unique router-id TLV.
(config-router)#mpls traffic-eng level-2	Enable MPLS-TE in is-type Level-2.
(config-router)#net 49.0001.0000.0000.0003.00	Configure the ISIS net address with area id as: 49.0001 and system-id as: 0000.0000.0003.
(config-router)#commit	Commit candidate configuration to the running configuration
(config-router)#exit	Exit Router mode.
(config)#interface eth2	Enter the interface configure mode.
(config-if)#bandwidth 1g	Configure Maximum link Bandwidth as 1g bits per second.
(config-if)#reservable-bandwidth 100m	Specify the maximum reservable bandwidth per interface.
(config-if)#commit	Commit candidate configuration to the running

© 2026 IP Infusion Inc. Proprietary

	configuration
(config-if)#exit	Exit interface configure mode.

Validation

R2

Check the output of show clns neighbors (page 2034) to verify that ISIS adjacency is up.

```
#show clns neighbors
Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 1
Total number of adjacencies: 3
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        5254.003d.2123  Up     6         L1   IS-IS
0000.0000.0003 eth2        5254.00fa.54d5  Up     6         L1   IS-IS
                                           Up     6         L2   IS-IS
```

Check the output of **show isis database (page 2056)** **level-1 verbose** to verify that LSP does have configured Max Link Bandwidth and Reservable Bandwidth.

```
R2#show isis database level-1 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  0x00000007   0xA4DF        1135              0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   13.1.1.1
  Router ID:    1.1.1.1
  Metric:   10         IS-Extended 0000.0000.0001.01
    IPv4 Interface Address: 13.1.1.1
    Neighbor IP Address: 13.1.1.1
    Maximum Link Bandwidth : 1000000.00 kbits/sec
    Reservable Bandwidth : 100000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 13.1.1.0/24
  Metric:   10         IP-Extended 1.1.1.1/32
0000.0000.0001.01-00  0x00000004   0x8B3C        683               0/0/0
  Metric:   0          IS-Extended 0000.0000.0001.00
  Metric:   0          IS-Extended 0000.0000.0002.00
0000.0000.0002.00-00* 0x0000000E   0x3961        1069              0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   13.1.1.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended 0000.0000.0001.01
    IPv4 Interface Address: 13.1.1.2
    Neighbor IP Address: 13.1.1.1
    Maximum Link Bandwidth : 500000.00 kbits/sec
    Reservable Bandwidth : 200000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 200000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 200000.00 kbits/sec
```

```
      Unreserved Bandwidth at priority 2: 200000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 200000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 200000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 200000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 200000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 200000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended 0000.0000.0003.02
    IPv4 Interface Address: 15.1.1.2
    Neighbor IP Address: 15.1.1.3
    Maximum Link Bandwidth : 1000000.00 kbits/sec
    Reservable Bandwidth : 100000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 2.2.2.0/24
  Metric:   10         IP-Extended 13.1.1.0/24
  Metric:   10         IP-Extended 15.1.1.0/24
0000.0000.0003.00-00  0x0000000A   0x4C9C        437               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   15.1.1.3
  Router ID:    3.3.3.3
  Metric:   10         IS-Extended 0000.0000.0003.02
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 15.1.1.0/24
0000.0000.0003.02-00  0x00000001   0x8C39        437               0/0/0
  Metric:   0          IS-Extended 0000.0000.0003.00
  Metric:   0          IS-Extended 0000.0000.0002.00
```

Check the output of **show isis database (page 2056)level-2 verbose** to verify that LSP does have configured Max Link Bandwidth and Reservable Bandwidth.

```
R2#show isis database level-2 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* 0x00000014   0x8C1A        1012              0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   15.1.1.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended 0000.0000.0003.02
    IPv4 Interface Address: 15.1.1.2
    Neighbor IP Address: 15.1.1.3
    Maximum Link Bandwidth : 1000000.00 kbits/sec
    Reservable Bandwidth : 100000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   20         IP-Extended 1.1.1.1/32
  Metric:   10         IP-Extended 15.1.1.0/24
```

```
  Metric:   10         IP-Extended 13.1.1.0/24
  Metric:   20         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 2.2.2.0/24
0000.0000.0003.00-00  0x0000000D   0x7B17        1006              0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   15.1.1.3
  Router ID:    3.3.3.3
  Metric:   10         IS-Extended 0000.0000.0003.02
    IPv4 Interface Address: 15.1.1.3
    Neighbor IP Address: 15.1.1.3
    Maximum Link Bandwidth : 1000000.00 kbits/sec
    Reservable Bandwidth : 100000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 15.1.1.0/24
  Metric:   30         IP-Extended 1.1.1.1/32
  Metric:   20         IP-Extended 13.1.1.0/24
  Metric:   20         IP-Extended 2.2.2.0/24
0000.0000.0003.02-00  0x00000001   0x8C39        375               0/0/0
  Metric:   0          IS-Extended 0000.0000.0003.00
  Metric:   0          IS-Extended 0000.0000.0002.00
```

R1

Check the output of show clns neighbors (page 2034) to verify that ISIS adjacency is up.

```
#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        5254.0028.3734      Up     24        L1   IS-IS
```

Check the output of **show isis database (page 2056) level-1 verbose** to verify that LSP does have configured Max Link Bandwidth and Reservable Bandwidth.

```
R1#show isis database level-1 verbose
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL    PDU Length
0000.0000.0001.00-00* 0x00000006   0xA6DE        331               0/0/0    143
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   13.1.1.1
  Router ID:    1.1.1.1
  Metric:   10         IS-Extended 0000.0000.0001.01
    IPv4 Interface Address: 13.1.1.1
    Neighbor IP Address: 13.1.1.1
    Maximum Link Bandwidth : 1000000.00 kbits/sec
    Reservable Bandwidth : 100000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000.00 kbits/sec
```

```
      Unreserved Bandwidth at priority 3: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 13.1.1.0/24
  Metric:   10         IP-Extended 1.1.1.1/32
0000.0000.0001.01-00* 0x00000004   0x8B3C        780              0/0/0    51
  Metric:   0          IS-Extended 0000.0000.0001.00
  Metric:   0          IS-Extended 0000.0000.0002.00
0000.0000.0002.00-00  0x0000000E   0x3961        1164             0/0/0    224
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   13.1.1.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended 0000.0000.0001.01
    IPv4 Interface Address: 13.1.1.2
    Neighbor IP Address: 13.1.1.1
    Maximum Link Bandwidth : 500000.00 kbits/sec
    Reservable Bandwidth : 200000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 200000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 200000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 200000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 200000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 200000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 200000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 200000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 200000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IS-Extended 0000.0000.0003.02
    IPv4 Interface Address: 15.1.1.2
    Neighbor IP Address: 15.1.1.3
    Maximum Link Bandwidth : 1000000.00 kbits/sec
    Reservable Bandwidth : 100000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 2.2.2.0/24
  Metric:   10         IP-Extended 13.1.1.0/24
  Metric:   10         IP-Extended 15.1.1.0/24
0000.0000.0003.00-00  0x0000000A   0x4C9C        533              0/0/0    80
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   15.1.1.3
  Router ID:    3.3.3.3
  Metric:   10         IS-Extended 0000.0000.0003.02
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 15.1.1.0/24
0000.0000.0003.02-00  0x00000001   0x8C39        533              0/0/0    51
  Metric:   0          IS-Extended 0000.0000.0003.00
  Metric:   0          IS-Extended 0000.0000.0002.00
```

R3

Check the output of show clns neighbors (page 2034) to verify that ISIS adjacency is up.

```
#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
0000.0000.0002 eth2        5254.00d4.2c13  Up     20        L1   IS-IS
                                           Up     20        L2   IS-IS
```

Check the output of **show isis database** (page 2056)**level-1 verbose** to verify that LSP does have configured Max Link Bandwidth and Reservable Bandwidth.

```
R3#show isis database level-2 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  0x00000010   0xA4E5        942               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   6.6.6.6
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended 0000.0000.0003.02
    IPv4 Interface Address: 15.1.1.2
    Neighbor IP Address: 15.1.1.3
    Maximum Link Bandwidth : 1000000.00 kbits/sec
    Reservable Bandwidth : 100000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 6.6.6.6/32
  Metric:   10         IP-Extended 15.1.1.0/24
  Metric:   10         IP-Extended 13.1.1.0/24
  Metric:   20         IP-Extended 1.1.1.1/32
  Metric:   20         IP-Extended 3.3.3.3/32
0000.0000.0003.00-00* 0x0000000B   0xF286        950               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   15.1.1.3
  Router ID:    3.3.3.3
  Metric:   10         IS-Extended 0000.0000.0003.02
    IPv4 Interface Address: 15.1.1.3
    Neighbor IP Address: 15.1.1.3
    Maximum Link Bandwidth : 1000000.00 kbits/sec
    Reservable Bandwidth : 100000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 100000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 100000.00 kbits/sec
    TE-Default Metric: 10
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 15.1.1.0/24
  Metric:   30         IP-Extended 1.1.1.1/32
  Metric:   20         IP-Extended 6.6.6.6/32
  Metric:   20         IP-Extended 13.1.1.0/24
0000.0000.0003.02-00* 0x00000001   0x8C39        939               0/0/0
  Metric:   0          IS-Extended 0000.0000.0003.00
```

```
Metric:   0         IS-Extended 0000.0000.0002.00
```

Administrative Group Constraints

To configure administrative group constraints (also known as color constraints) in Level-1 and Level-2 on L1-L2 IS:

- Configure support for required admin groups in NSM on all participating routers
- Configure required administrative groups on all participating interfaces

The configuration in this example forces the primary LSP to be setup through links that belong either to admin group A or C. A link that does not belong to either of these admin groups will not be used for setting up the LSP.

Figure 139. MPLS-TE Topology

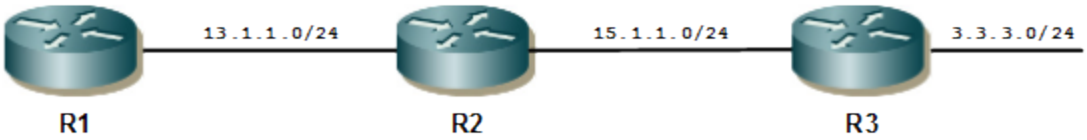

Configuration

R2

`#configure terminal`	Enter Configure mode.
`(config)#interface eth1`	Specify the interface (`eth1`) to configure, and enter Interface mode.
`(config-if)#ip address 13.1.1.2/24`	Configure ip address to the interface eth1.
`(config-if)#ip router isis 1`	Enable ISIS on interface eth1.
`(config-if)#isis circuit-type level-1`	Configure ISIS circuit-type as Level-1
`(config)#interface eth2`	Enter Interface eth2 configure mode.
`(config-if)#ip address 15.1.1.2/24`	Configure ip address to the interface eth2.
`(config-if)#ip router isis 1`	Enable ISIS on interface eth2.
`(config-if)#isis circuit-type level-2`	Configure ISIS circuit-type as Level-2.
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-if)#exit`	Exit from interface configuration mode.
`(config)#admin-group a 0`	Add new administrative groups and specify their names and assign bit values to them.
`(config)#admin-group b 1`	
`(config)#admin-group c 2`	
`(config)#admin-group d 3`	
`(config)#router isis 1`	Create an IS-IS routing instance for area 49 (1).
`(config-router)#metric-style wide`	Configure the new style of metric type as wide.
`(config-router)#mpls traffic-eng router-id 2.2.2.2`	Configure MPLS-TE unique router-id TLV.

© 2026 IP Infusion Inc. Proprietary

`(config-router)#mpls traffic-eng level-1`	Enable MPLS-TE in is-type Level-1.
`(config-router)#mpls traffic-eng level-2`	Enable MPLS-TE in is-type Level-2.
`(config-router)#net 49.0001.0000.0000.0002.00`	Configure the ISIS net address with area id as: 49.0001 and system id as: 0000.0000.0002.
`(config-router)#commit`	Commit candidate configuration to the running configuration
`(config-router)#exit`	Exit Router mode.
`(config)#interface eth1`	Enter the interface configure mode.
`(config-if)#admin-group a`	Add administrative groups to the links. When used in the interface mode, this command adds a link between an Interface and a group. The name is the name of the group Previously configured. You can have multiple groups per Interface.
`(config-if)#admin-group c`	
`(config-if)#exit`	Exit interface configure mode.
`(config)#interface eth2`	Enter the interface configure mode.
`(config-if)#admin-group b`	Add administrative groups to the links. When used in the interface mode, this command adds a link between an Interface and a group. The name is the name of the group previously configured. You can have multiple groups per Interface.
`(config-if)#admin-group d`	
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-if)#end`	Exit interface configure mode.

R1

`#configure terminal`	Enter Configure mode.
`(config)#interface eth1`	Specify the interface (`eth1`) to configure, and enter Interface mode.
`(config-if)#ip address 13.1.1.1/24`	Configure ip address to the interface eth1.
`(config-if)#ip router isis 1`	Enable ISIS on interface eth1.
`(config-if)#isis circuit-type level-1`	Configure ISIS circuit-type as Level-1
`(config-if)#commit`	Commit candidate configuration to the running configuration
`(config-if)#exit`	Exit interface configure mode.
`(config)#admin-group a 0`	Add new administrative groups and specify their names and assign bit values to them.
`(config)#admin-group c 2`	
`(config)#router isis 1`	Create an IS-IS routing instance for area 49 (1).
`(config-router)#metric-style wide`	Configure the new style of metric type as wide.
`(config-router)#is-type level-1`	Configure IS-Type as Level-1 router.

© 2026 IP Infusion Inc. Proprietary

(config-router)#mpls traffic-eng router-id 1.1.1.1	Configure MPLS-TE unique router-id TLV.
(config-router)#mpls traffic-eng level-1	Enable MPLS-TE in is-type Level-1.
(config-router)#net 49.0001.0000.0000.0001.00	Configure the ISIS net address with area id as: 49.0001 and system-id as: 0000.0000.0001.
(config-router)#commit	Commit candidate configuration to the running configuration
(config-router)#exit	Exit Router mode.
(config)#interface eth1	Enter the interface configure mode.
(config-if)#admin-group a	Add administrative groups to the links. When used in the interface mode, this command adds a link between an Interface and a group. The name is the name of the group Previously configured. You can have multiple groups per Interface.
(config-if)#admin-group c	
(config-if)#commit	Commit candidate configuration to the running configuration
(config-if)#end	Exit interface configure mode.

R3

#configure terminal	Enter Configure mode.
(config)#interface eth2	Specify the interface (eth2) to configure, and enter Interface mode.
(config-if)#ip address 15.1.1.3/24	Configure ip address to the interface eth1.
(config-if)#ip router isis 1	Enable ISIS on interface eth1.
(config-if)#isis circuit-type level-2	Configure ISIS circuit-type as Level-1
(config-if)#commit	Commit candidate configuration to the running configuration
(config-if)#exit	Exit interface configure mode.
(config)#admin-group b 1	Add new administrative groups and specify their names and assign bit values to them.
(config)#admin-group d 3	
(config)#router isis 1	Create an IS-IS routing instance for area 49 (1).
(config-router)#metric-style wide	Configure the new style of metric type as wide.
(config-router)#is-type level-2	Configure IS-Type as Level-2 router.
(config-router)#mpls traffic-eng router-id 3.3.3.3	Configure MPLS-TE unique router-id TLV.
(config-router)#mpls traffic-eng level-2	Enable MPLS-TE in is-type Level-2.
(config-router)#net 49.0001.0000.0000.0003.00	Configure the ISIS net address with area id as:

	49.0001 and system-id as: 0000.0000.0003.
(config-router)#commit	Commit candidate configuration to the running configuration
(config-router)#exit	Exit Router mode.
(config)#interface eth2	Enter the interface configure mode.
(config-if)#admin-group b	Add administrative groups to the links. When used in the interface mode, this command adds a link between an Interface and a group. The name is the name of the group
(config-if)#admin-group d	Previously configured. You can have multiple groups per Interface.
(config-if)#commit	Commit candidate configuration to the running configuration
(config-if)#end	Exit interface configure mode.

Validation

R2

Check the output of show clns neighbors (page 2034) to verify that ISIS adjacency is up.

```
#show clns neighbors

Area 1:
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        5254.00bb.5e85      Up     24        L1   IS-IS
0000.0000.0003 eth2        5254.00ac.f960      Up     9         L2   IS-IS
```

Check the output of **show isis database (page 2056)level-1 verbose** to verify that LSP does have configured admin-group constraints in Level-1.

```
  straints in Level-1.
R2#show isis database level-1 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   0x3022        1020              0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   13.1.1.1
  Router ID:    1.1.1.1
  Metric:   10         IS-Extended 0000.0000.0001.01
    Admin-Group:
      Group 0
      Group 2
    IPv4 Interface Address: 13.1.1.1
    Neighbor IP Address: 13.1.1.1
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
```

```
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 10
  Metric:   10         IP-Extended 13.1.1.0/24
0000.0000.0001.01-00  0x00000001   0x8F3D        820              0/0/0
  Metric:   0          IS-Extended 0000.0000.0001.00
  Metric:   0          IS-Extended 0000.0000.0002.00
0000.0000.0002.00-00* 0x00000004   0xBB8D        1070             0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   13.1.1.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended 0000.0000.0001.01
    Admin-Group:
      Group 0
      Group 2
    IPv4 Interface Address: 13.1.1.2
    Neighbor IP Address: 13.1.1.1
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 10
  Metric:   10         IP-Extended 13.1.1.0/24
```

Check the output of **show isis database (page 2056) level-2 verbose** to verify that LSP does have configured admin-group constraints in Level-2.

```
R2#show isis database level-2 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   0x1FDE        1058              0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   15.1.1.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended 0000.0000.0003.01
    Admin-Group:
      Group 1
      Group 3
    IPv4 Interface Address: 15.1.1.2
    Neighbor IP Address: 15.1.1.3
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 10
```

```
  Metric:   10         IP-Extended 15.1.1.0/24
  Metric:   10         IP-Extended 13.1.1.0/24
0000.0000.0003.00-00  0x00000004   0x71BF        1081              0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   15.1.1.3
  Router ID:    3.3.3.3
  Metric:   10         IS-Extended 0000.0000.0003.01
    Admin-Group:
      Group 1
      Group 3
    IPv4 Interface Address: 15.1.1.3
    Neighbor IP Address: 15.1.1.3
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 10
  Metric:   10         IP-Extended 15.1.1.0/24
0000.0000.0003.01-00  0x00000001   0x9333        827               0/0/0
  Metric:   0          IS-Extended 0000.0000.0003.00
  Metric:   0          IS-Extended 0000.0000.0002.00
```

R1

Check the output of show clns neighbors (page 2034) to verify that ISIS adjacency is up.

```
#show clns neighbors
Area 1:
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        5254.00f6.4ae7      Up     7         L1   IS-IS
```

Check the output of **show isis database (page 2056)level-1 verbose** to verify that LSP does have configured admin-group constraints in Level-1.

```
R1#show isis database level-1 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   0x3022        1105              0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   13.1.1.1
  Router ID:    1.1.1.1
  Metric:   10         IS-Extended 0000.0000.0001.01
    Admin-Group:
      Group 0
      Group 2
    IPv4 Interface Address: 13.1.1.1
    Neighbor IP Address: 13.1.1.1
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
```

```
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 10
  Metric:   10         IP-Extended 13.1.1.0/24
0000.0000.0001.01-00* 0x00000001   0x8F3D        905               0/0/0
  Metric:   0          IS-Extended 0000.0000.0001.00
  Metric:   0          IS-Extended 0000.0000.0002.00
0000.0000.0002.00-00  0x00000004   0xBB8D        1153              0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   13.1.1.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended 0000.0000.0001.01
    Admin-Group:
      Group 0
      Group 2
    IPv4 Interface Address: 13.1.1.2
    Neighbor IP Address: 13.1.1.1
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 10
  Metric:   10         IP-Extended 13.1.1.0/24
```

R3

Check the output of show clns neighbors (page 2034) to verify that ISIS adjacency is up.

```
#show clns neighbors
Area 1:
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth2        5254.0062.3ea6      Up     28        L2   IS-IS
```

Check the output of **show isis database (page 2056) level-2 verbose** to verify that LSP does have configured admin-group constraints in Level-2.

```
R3#show isis database level-2 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   0x1FDE        1037              0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   15.1.1.2
  Router ID:    2.2.2.2
  Metric:   10         IS-Extended 0000.0000.0003.01
    Admin-Group:
      Group 1
      Group 3
    IPv4 Interface Address: 15.1.1.2
    Neighbor IP Address: 15.1.1.3
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
      Unreserved Bandwidth at priority 3: 10g
```

```
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 10
  Metric:   10         IP-Extended 15.1.1.0/24
  Metric:   10         IP-Extended 13.1.1.0/24
0000.0000.0003.00-00* 0x00000004   0x71BF        1063              0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  IP Address:   15.1.1.3
  Router ID:    3.3.3.3
  Metric:   10         IS-Extended 0000.0000.0003.01
    Admin-Group:
      Group 1
      Group 3
    IPv4 Interface Address: 15.1.1.3
    Neighbor IP Address: 15.1.1.3
    Maximum Link Bandwidth: 10g
    Reservable Bandwidth: 10g
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10g
      Unreserved Bandwidth at priority 1: 10g
      Unreserved Bandwidth at priority 2: 10g
      Unreserved Bandwidth at priority 3: 10g
      Unreserved Bandwidth at priority 4: 10g
      Unreserved Bandwidth at priority 5: 10g
      Unreserved Bandwidth at priority 6: 10g
      Unreserved Bandwidth at priority 7: 10g
    TE-Default Metric: 10
  Metric:   10         IP-Extended 15.1.1.0/24
0000.0000.0003.01-00* 0x00000001   0x9333        808               0/0/0
  Metric:   0          IS-Extended 0000.0000.0003.00
  Metric:   0          IS-Extended 0000.0000.0002.00
```

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:

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

```
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 140. ISIS Multi Topology

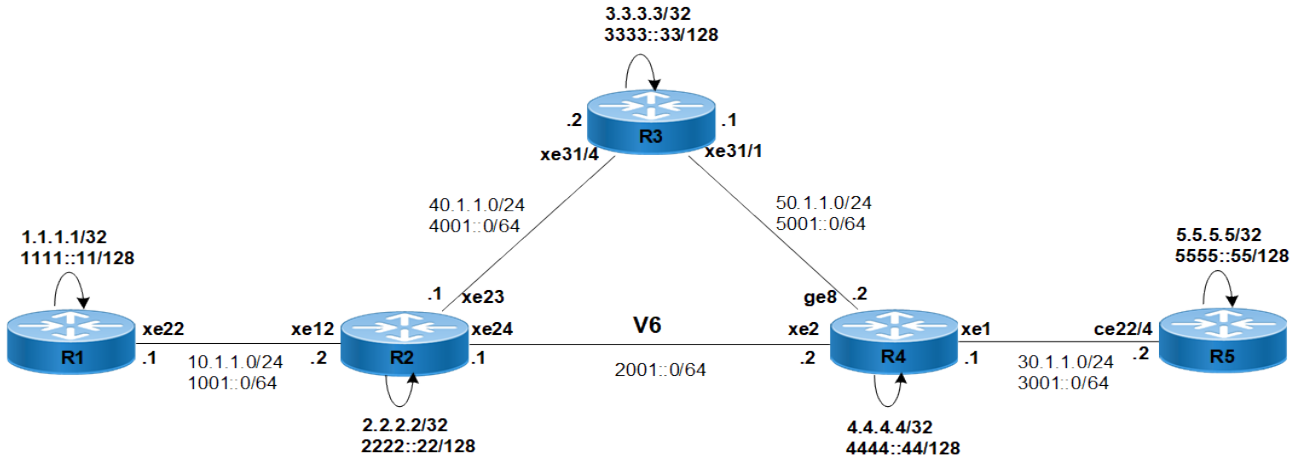

To configure multi topology on the routers R1, R2, R3, R4 and R5, follow the steps mentioned below:

Notes:

- Ensure that the Prerequisites (page 1895) are met for all the routers.
- 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

```
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
```

© 2026 IP Infusion Inc. Proprietary

```
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
```

© 2026 IP Infusion Inc. Proprietary

```
  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
```

© 2026 IP Infusion Inc. Proprietary

```
  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
```

© 2026 IP Infusion Inc. Proprietary

```
  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
```

© 2026 IP Infusion Inc. Proprietary

```
  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.

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

INTERMEDIATE SYSTEM TO INTERMEDIATE SYSTEM COMMAND REFERENCE

IS-IS Commands 1929
- accept-lifetime 1932
- address-family ipv6 1934
- adjacency-check 1935
- area-password 1936
- authentication key-chain 1937
- authentication mode 1938
- authentication send-only 1939
- bfd all-interfaces 1940
- capability cspf 1941
- clear clns is-neighbors 1942
- clear clns neighbors 1943
- clear ip isis route 1944
- clear isis adjacency 1945
- clear isis counter 1946
- clear isis interface counter 1947
- clear isis process 1948
- debug isis 1949
- default-information originate 1951
- distance (IPv4) 1952
- distance (IPv6) 1954
- domain password 1956
- dynamic-hostname 1958
- fast-reroute per-prefix 1959
- fast-reroute per-prefix remote-lfa 1960
- fast-reroute terminate-hold-on interval 1961
- fast-reroute tie-break 1962
- ignore-lsp-errors 1964
- ip route high-priority tag 1965
- ip router isis 1966
- ipv6 router isis 1967
- is-type 1968
- isis authentication key-chain 1969
- isis authentication mode md5 1970
- isis authentication send-only 1971

© 2026 IP Infusion Inc. Proprietary

isis bfd 1972
isis circuit-type 1973
isis csnp-interval 1974
isis fast-reroute per-prefix candidate disable 1975
isis hello-interval 1976
isis hello-multiplier 1978
isis hello padding 1979
isis lsp-interval 1980
isis mesh-group 1981
isis metric 1982
isis network 1983
isis password 1984
isis priority 1985
isis retransmit-interval 1986
isis tag 1987
isis te-metric 1988
isis wait-timer 1989
isis wide-metric 1990
ispf 1991
key chain 1992
key-id 1993
key-string 1994
key-string encrypted 1995
lsp-gen-interval 1996
lsp-mtu 1997
lsp-refresh-interval 1998
max-area-address 1999
max-lsp-lifetime 2000
metric-style 2001
mpls traffic-eng 2003
mpls traffic-eng router-id 2004
net 2005
passive-interface 2006
prc-interval-exp 2007
redistribute 2008
redistribute isis 2010
redistribute isis WORD 2012
router isis 2014
send-lifetime 2015
set-overload-bit 2017

© 2026 IP Infusion Inc. Proprietary

snmp context ... 2019
spf-interval-exp ... 2020
summary-address ... 2021
summary-prefix ... 2022
IS-IS Graceful Restart Commands ... 2023
capability restart graceful ... 2024
isis restart grace-period ... 2025
isis restart-hello-interval ... 2026
isis restart helper ... 2027
isis restart suppress-adjacency ... 2028
restart isis graceful ... 2029
restart-timer ... 2030
ISIS Show Commands ... 2031
show clns is-neighbors ... 2032
show clns neighbors ... 2034
show cspf lsp ... 2035
show cspf rsvp forwarding-timer ... 2038
show debugging isis ... 2039
show ip isis igp-shortcut-lsp ... 2040
show ip isis lfa-config ... 2041
show ip isis route ... 2042
show ip isis route fast-reroute ... 2044
show ip isis route igp-shortcut ... 2046
show ip protocols ... 2048
show ip route fast-reroute ... 2050
show ipv6 isis route ... 2051
show ipv6 isis topology ... 2053
show isis counter ... 2055
show isis database ... 2056
show isis interface ... 2059
show isis microloop-avoidance ... 2062
show isis spf-logs ... 2064
show isis topology ... 2065
show running-config interface isis ... 2067
show running-config router isis ... 2068
IS-IS Microloop-Avoidance Commands ... 2069
microloop-avoidance ... 2070
microloop-avoidance flex-algo ... 2072
microloop-avoidance hold-timer ... 2073
microloop-avoidance max-fib ... 2074

© 2026 IP Infusion Inc. Proprietary

microloop-avoidance segment-routing 2075
show isis microloop-avoidance 2076

IS-IS Commands

This section describes the IS-IS commands:

accept-lifetime 1932
address-family ipv6 1934
adjacency-check 1935
area-password 1936
authentication key-chain 1937
authentication mode 1938
authentication send-only 1939
bfd all-interfaces 1940
capability cspf 1941
clear clns is-neighbors 1942
clear clns neighbors 1943
clear ip isis route 1944
clear isis adjacency 1945
clear isis counter 1946
clear isis interface counter 1947
clear isis process 1948
debug isis 1949
default-information originate 1951
distance (IPv4) 1952
distance (IPv6) 1954
domain password 1956
dynamic-hostname 1958
fast-reroute per-prefix 1959
fast-reroute per-prefix remote-lfa 1960
fast-reroute terminate-hold-on interval 1961
fast-reroute tie-break 1962
ignore-lsp-errors 1964
ip route high-priority tag 1965
ip router isis 1966
ipv6 router isis 1967
is-type 1968
isis authentication key-chain 1969
isis authentication mode md5 1970
isis authentication send-only 1971
isis bfd 1972
isis circuit-type 1973

© 2026 IP Infusion Inc. Proprietary

isis csnp-interval 1974
isis fast-reroute per-prefix candidate disable 1975
isis hello-interval 1976
isis hello-multiplier 1978
isis hello padding 1979
isis lsp-interval 1980
isis mesh-group 1981
isis metric 1982
isis network 1983
isis password 1984
isis priority 1985
isis retransmit-interval 1986
isis tag 1987
isis te-metric 1988
isis wait-timer 1989
isis wide-metric 1990
ispf 1991
key chain 1992
key-id 1993
key-string 1994
key-string encrypted 1995
lsp-gen-interval 1996
lsp-mtu 1997
lsp-refresh-interval 1998
max-area-address 1999
max-lsp-lifetime 2000
metric-style 2001
mpls traffic-eng 2003
mpls traffic-eng router-id 2004
net 2005
passive-interface 2006
prc-interval-exp 2007
redistribute 2008
redistribute isis 2010
redistribute isis WORD 2012
router isis 2014
send-lifetime 2015
set-overload-bit 2017
snmp context 2019
spf-interval-exp 2020

© 2026 IP Infusion Inc. Proprietary

summary-address 2021
summary-prefix 2022

© 2026 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 disable it.

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 expire.

duration
Specify the duration of the key in seconds

<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

Key-chain key mode

© 2026 IP Infusion Inc. Proprietary

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the setting of accept-lifetime for key-id 1 on the key chain named mychain.

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key-id 1
(config-keychain-key)#accept-lifetime 03:03:01 Dec 03 2004 04:04:02 Oct 06 2006
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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 (RSVP-TE) 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
```

© 2026 IP Infusion Inc. Proprietary

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

Execution modeand Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
>enable
#clear clns is-neighbors 1111.1111.1111.1111
```

© 2026 IP Infusion Inc. Proprietary

clear clns neighbors

Use this command to clear CLNS neighbor adjacencies.

Command Syntax

```
clear clns neighbors
```

Parameters

None

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
>ena
#clear clns neighbors
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
>ena
#clear ip isis route redistribution
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear isis adjacency *
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear isis counter
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear isis interface counter
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear isis process
```

© 2026 IP Infusion Inc. Proprietary

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 (ifsm|nfsm|events|pdu|lsp|ofib|sr-micro-loop-
avoidance|spf|nsm|rib|checksum|authentication|protocol-errors|local-
updates|bfd|mpls|sr|lfa|tilfa|asla|flexalgo|dist-ls)
no debug isis (all|)
no debug all
no debug all isis
no debug isis (ifsm|nfsm|events|pdu|lsp|ofib|sr-micro-loop-
avoidance|spf|nsm|rib|checksum|authentication|protocol-errors|local-
updates|bfd|mpls|sr|lfa|tilfa|asla|flexalgo|dist-ls)
```

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.

mpls
Debugging for MPLS.

IFNAME
Interface name.

System-ID
System identifier.

cspf
Debugging for CSPF.

events
Events

hexdump
Hexdump.

ifsm
Debugging for interface finite state machine.

local-updates
Debugging for local updates.

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#debug isis pdu

#configure terminal
(config)#debug isis nsm
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

distance (IPv4)

Use the distance command in router mode to configure the administrative distance for:

- All IPv4 routes received from a specific source System-ID
- Optionally, routes that match a specified access list

Use the `no` form:

- To remove a previously configured administrative distance, use the no distance command.
- To remove distance settings configured with a specific System-ID or access-list use no distance along with system-id only.

Command Syntax

```
distance <1-255> (System-ID (WORD|)|)
no distance
no distance (System-ID|)
```

Parameters

<1-255>
Distance range.

System-ID
Source ID in XXXX.XXXX.XXXX format.

WORD
Access-list name.

Default

Turned off.

Command Mode

Router 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)#distance 10
```

The following example shows setting the administrative distance to 40, only for routes learned from the router with System-ID 0000.0000.0001

```
#configure terminal
(config)#router isis
(config-router)#distance 40 0000.0000.0001
```

© 2026 IP Infusion Inc. Proprietary

The following example shows setting the administrative distance to 40, only for routes from the system with System-id that match access list DIST.

```
#configure terminal
(config)#router isis
(config-router)#distance 40 0000.0000.0001 DIST
```

© 2026 IP Infusion Inc. Proprietary

distance (IPv6)

Use the distance command in router mode to configure the administrative distance for:

- All IPv6 routes received from a specific source System-ID
- Optionally, routes that match a specified access list

Use the `no` form:

- To remove a previously configured administrative distance,
- To remove distance settings configured with a specific System-ID or access-list use no distance along with system-id only.

Command Syntax

```
distance <1-255>
no distance
```

Parameters

<1-255>
Distance range.

Default

Turned off

Command Mode

Address-family ipv6 mode

Applicability

This command was introduced before OcNOS version 1.3. T

Examples

The following example shows setting the administrative distance for all ipv6 routes.

```
#configure terminal
(config)#router isis
(config-router)address-family ipv6
(config-router-af)#distance 10
```

The following example shows setting the administrative distance to 40 but only for routes learned from the router with System-ID 0000.0000.0001

```
#configure terminal
(config)#router isis
(config-router)address-family ipv6
(config-router-af)distance 40 0000.0000.0001
```

The following example shows setting the administrative distance to 40 but only for routes from the system with System-id that match access list DIST

© 2026 IP Infusion Inc. Proprietary

```
#configure terminal
(config)#router isis
(config-router)address-family ipv6
(config-router-af)distance 40 0000.0000.0001 DIST
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

```
(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

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
```

© 2026 IP Infusion Inc. Proprietary

fast-reroute per-prefix remote-lfa

Use this command to configure remote loop-free alternate fast reroute configuration for a particular level of an IS-IS instance.

Use the `no` parameter to remove remote loop-free alternate fast reroute configuration.

Command Syntax

```
fast-reroute per-prefix remote-lfa (level-1|level-2) proto ipv4 tunnel mpls-ldp
no fast-reroute per-prefix remote-lfa (level-1|level-2) proto ipv4
```

Parameters

level-1
Level 1 only.

level-2
Level 2 only.

Default

Disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 5.1.

Examples

```
#configure terminal
(config)#router isis 1
(config-router)#fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp
(config-router)#
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

Disabled

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
```

© 2026 IP Infusion Inc. Proprietary

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

Disabled

Command Mode

Interface mode

Example

```
#configure terminal
(config)#interface eth0
(config-if)#ipv6 router isis bb
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

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.

© 2026 IP Infusion Inc. Proprietary

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#isis hello-interval 5 level-1

(config-if)#isis hello-interval minimal level-1-only
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

Disabled

Command Mode

Interface mode

Examples

```
(config)#interface eth0
(config-if)#isis mesh-group 20

(config)#interface eth2
(config-if)#no isis mesh-group
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

isis te-metric

Use this command to set 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 TE 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

ISIS wide metric value.

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS version 5.0

Examples

```
#configure terminal
(config)#interface eth1
(config-router)#isis te-metric 100
(config)#interface eth1
(config-router)#no isis te-metric
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

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)#
```

© 2026 IP Infusion Inc. Proprietary

key-id

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-id <0-2147483647>
no key-id <0-2147483647>
```

Parameters

<0-2147483647>
Specify a key identifier.

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

Key-chain 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-id 1
(config-keychain-key)#
```

© 2026 IP Infusion Inc. Proprietary

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

Key-chain mode and Key-chain key mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

In the following example, the password for `key-id 1` in the key chain named `mychain` is set to `prime`:

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key-id 1
(config-keychain-key)#key-string prime

(config-keychain)#key-id 1
(config-keychain-key)#no key-string
```

© 2026 IP Infusion Inc. Proprietary

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

Key-chain mode and Key-chain key mode

Applicability

This command was introduced in OcNOS version 4.1.

Example

In the following example, the encrypted password for key-id 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-id 1
(config-keychain-key)#key-string encrypted 0xd6c50b442de47f70
(config-keychain)#key-id 1
(config-keychain-key)#no key-string
```

© 2026 IP Infusion Inc. Proprietary

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

ISIS uses 5 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 10
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 2000) 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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 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 (page 1998) 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

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
```

© 2026 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

ISIS uses narrow metric style for level 1 and 2.

© 2026 IP Infusion Inc. Proprietary

Command Mode

Router mode

Examples

```
(config)#router isis bb
(config-router)#metric-style narrow transition
```

© 2026 IP Infusion Inc. Proprietary

mpls traffic-eng

Use this command to configure MPLS Traffic Engineering feature for ISIS.

Use the `no` parameter to turn off the feature.

Note: Metric-style wide must be enabled before using this CLI.

Command Syntax

```
mpls traffic-eng (level-1|level-2)
no mpls traffic-eng (level-1|level-2)
```

Parameters

level-1

Specify the level-1 MPLS Traffic Engineering feature.

level-2

Specify the level-2 MPLS Traffic Engineering feature.

Default

If this command is not used, ISIS does not encode traffic engineering TLVs and Sub-TLVs.

Command Mode

Router mode

Examples

```
(config)#router isis bb
(config-router)#metric-style wide
(config-router)#mpls traffic-eng level-1

(config)#router isis bb
(config-router)#no mpls traffic-eng level-1
```

© 2026 IP Infusion Inc. Proprietary

mpls traffic-eng router-id

Use this command to configure the traffic engineering stable IP address for a system.

Use the `no` parameter to turn off the feature.

Command Syntax

```
mpls traffic-eng router-id A.B.C.D
no mpls traffic-eng router-id
```

Parameters

A.B.C.D

Specify the ISIS router-ID in an IP address format.

Default

If this command is not used, and traffic engineering is enabled, ISIS will use global router-id..

Command Mode

Router mode

Examples

```
(config)#router isis bb
(config-router)#mpls traffic-eng router-id 10.10.0.23

(config)#router isis bb
(config-router)#no mpls traffic-eng router-id
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

redistribute

Use this command to redistribute routes from another protocol into the ISIS routing table.

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

Command Syntax

```
redistribute (kernel|connected|static|rip|ospf|bgp) {metric <0-16777215>| metric-type
(internal|external)|level-1|level-2|level-1-2|route-map WORD}
no redistribute (kernel|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
Metric for redistributed routes.

<0-16777215>
IS-IS default metric.

metric-type
IS-IS exterior metric type for redistributed routes:

internal
IS-IS internal metric type.

external
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
Route map reference.

WORD
Route map name.

© 2026 IP Infusion Inc. Proprietary

Default

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

```
(config)#router isis A
(config-router)#redistribute bgp metric 12
```

© 2026 IP Infusion Inc. Proprietary

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` form of this command to disable 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

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
```

© 2026 IP Infusion Inc. Proprietary

```
(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
```

© 2026 IP Infusion Inc. Proprietary

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

Disabled

Command Mode

Router mode

© 2026 IP Infusion Inc. Proprietary

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#router isis A
(config-router)#redistribute isis B metric 12
```

© 2026 IP Infusion Inc. Proprietary

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 (page 1966) 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

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)#
```

© 2026 IP Infusion Inc. Proprietary

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.

Default

No default value is specified

Command Mode

Keychain-key mode

© 2026 IP Infusion Inc. Proprietary

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the setting of `send-lifetime` for `key-id 1` on the key chain named `mychain`:

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key-id 1
(config-keychain-key)#send-lifetime 03:03:01 Jan 03 2004 04:04:02 Dec 06 2006
```

© 2026 IP Infusion Inc. Proprietary

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

Not set

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.

© 2026 IP Infusion Inc. Proprietary

```
OcNOS(config)#router isis bb
OcNOS(config-router)#set-overload-bit on-startup wait-for-bgp suppress interlevel external
```

© 2026 IP Infusion Inc. Proprietary

snmp context

Use this command to configure the SNMP context name that is mapped to each isis instance.

When context name is configured isis OID is registered with SNMP to that context.

Use no form of the command to un-map the SNMP context name from isis instance and unregister isis OID from SNMP.

Command Syntax

```
snmp context-name WORD
no snmp context-name
```

Parameters

WORD

SNMP context name (Maximum 32 alphanumeric characters)

Default

Disabled

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#router isis 1
OcNOS(config-router)#snmp context-name context1
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

IS-IS Graceful Restart Commands

This section describes the IS-IS graceful restart commands:

capability restart graceful 2024
isis restart grace-period 2025
isis restart-hello-interval 2026
isis restart helper 2027
isis restart suppress-adjacency 2028
restart isis graceful 2029
restart-timer 2030

© 2026 IP Infusion Inc. Proprietary

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

None

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 execution 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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

ISIS Show Commands

This section describes the ISIS Show commands.

show clns is-neighbors 2032
show clns neighbors 2034
show cspf lsp 2035
show cspf rsvp forwarding-timer 2038
show debugging isis 2039
show ip isis igp-shortcut-lsp 2040
show ip isis lfa-config 2041
show ip isis route 2042
show ip isis route fast-reroute 2044
show ip isis route igp-shortcut 2046
show ip protocols 2048
show ip route fast-reroute 2050
show ipv6 isis route 2051
show ipv6 isis topology 2053
show isis counter 2055
show isis database 2056
show isis interface 2059
show isis microloop-avoidance 2062
show isis spf-logs 2064
show isis topology 2065
show running-config interface isis 2067
show running-config router isis 2068

© 2026 IP Infusion Inc. Proprietary

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

Privileged execution 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 50. 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.
State	Adjacency state:

Table 50. show clns is-neighbors output (continued)

Field	Description
	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.

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution 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 51. 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.

© 2026 IP Infusion Inc. Proprietary

Table 51. 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 cspf lsp

Use this command to display CSPF statistics, including counts or status of CSPF-computed paths for label-switched paths. Displays the actual CSPF-selected paths with extensive constraint detail.

Command Syntax

```
Show cspf lsp
Show cspf lsp SESSION-NAME
show cspf lsp count
```

Parameters

SESSION-NAME

Use this parameter to display only the CSPF LSPs matching a specific RSVP session name.

count

Use this parameter to display the LSP count.

Command Mode

Privileged execution mode and Execution mode

Applicability

This command was introduced before OcNOS version 7.0.0.

Example

Following is a sample output from the command.

```
OcNOS#show cspf lsp
   LSP Name        : t1-Primary
   Trunk ID        : 5001
   Lsp Id          : 2202
   OSPF ID         : 1
   Ingress         : 2.2.2.2
   Egress          : 3.3.3.3
   Ext Tunnel ID   : 2.2.2.2
   LSP Type        : 0
   Client ID       : 1
   State           : Established(3)
   Setup Priority  : 7
   Hold Priority   : 0
   Hop Limit       : 255
   Include Mask    : 0x0
   Exclude Mask    : 0x0
   Include All Mask: 0x0
   LSP Metric      : 1
   LSP Reoptimize  : Yes
   LSP in ReopList : Yes
   LSP ReopTime    : 1 minutes (Offset 18 seconds)
   Next Reop Due in: 60 seconds
   Computed ERO    :
       30.0.0.3

   LSP Name        : BL-30.0.0.3-3.3.3.3-100-Bypass
   Trunk ID        : 5002
   Lsp Id          : 2201
   OSPF ID         : 1
   Ingress         : 2.2.2.2
   Egress          : 3.3.3.3
   Ext Tunnel ID   : 2.2.2.2
   LSP Type        : 2
   Client ID       : 1
   State           : Established(3)
   Setup Priority  : 7
   Hold Priority   : 0
   Hop Limit       : 255
   Include Mask    : 0x0
   Exclude Mask    : 0x0
   Include All Mask: 0x0
   LSP Metric      : 1
   LSP Reoptimize  : Yes
   LSP in ReopList : Yes
   LSP ReopTime    : 1 minutes (Offset 13 seconds)
   Next Reop Due in: 55 seconds
   Exclude Path Constraint :
       30.0.0.3    exclude link
   Computed ERO    :
       37.0.0.3

OcNOS#
```

© 2026 IP Infusion Inc. Proprietary

Following is a sample output from the command using the `SESSION-NAME` parameter.

```
OcNOS#show cspf lsp BL
 LSP detail
   LSP Name               : BL-30.0.0.3-3.3.3.3-100-Bypass
   Trunk ID               : 5002
   LSP ID                 : 2202
   ISIS Tag Name          : 10
   Client ID              : 5
   State                  : Established(3)
   Ingress                : 2.2.2.2
   Egress                 : 3.3.3.3
   Ext Tunnel ID          : 2.2.2.2
   LSP Type               : 2
   Setup Priority         : 7
   Hold Priority          : 0
   Hop Limit              : 255
   Include Mask           : 0x0
   Exclude Mask           : 0x0
   Include All Mask       : 0x0
   LSP Metric             : 10
   LSP Reoptimize         : Yes
   LSP in Reop-List       : Yes
   LSP ReopTime           : 1 minutes (Offset 14 seconds)
   Next Reop Due in       : 23 seconds
   Computed ERO           :
       37.0.0.3

OcNOS#
```

Following is a sample output from the command using the `count` parameter.

```
OcNOS#show cspf lsp count
 ISIS cspf lsp count
  LSP Count for Setup priority 0     : 0
  LSP Count for Setup priority 1     : 0
  LSP Count for Setup priority 2     : 0
  LSP Count for Setup priority 3     : 0
  LSP Count for Setup priority 4     : 0
  LSP Count for Setup priority 5     : 0
  LSP Count for Setup priority 6     : 0
  LSP Count for Setup priority 7     : 2
OcNOS#
```

show cspf rsvp forwarding-timer

This command displays the information of Graceful Restart capable RSVP client to ISIS or OSPF, CSPF that are currently shutdown.

Command Syntax

```
show cspf rsvp forwarding-timer
```

Parameters

None

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 5.0.

Example

```
#show cspf rsvp forwarding-timer
CSPF Server      Protocol-Name   GR-State        Time Remaining (sec)    Disconnected-time
OSPF             RSVP            ACTIVE          88                      2021/08/18 04:49:23
```

© 2026 IP Infusion Inc. Proprietary

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

Execution modePrivileged execution 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
```

© 2026 IP Infusion Inc. Proprietary

show ip isis igp-shortcut-lsp

Use this command to display IS-IS shortcut MPLS label-switched paths (LSPs).

Command Syntax

```
show ip isis (WORD|) igp-shortcut-lsp
```

Parameters

WORD
Information for a single IS-IS area.

Command Mode

Execution modeand Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
# show ip isis igp-shortcut-lsp
```

Table 52. 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.

© 2026 IP Infusion Inc. Proprietary

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

Execution modeand Privileged execution mode

Default

None

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
```

© 2026 IP Infusion Inc. Proprietary

show ip isis route

Use this command to display IS-IS routing table for IPv4.

Command Syntax

```
show ip isis route (algorithm|detail|fast-reroute|igp-shortcut|prefix|tilfa)
```

Parameters

word
Information for an IS-IS area.

algorithm
Displays information of the flexible algorithm.

detail
Displays detailed information.

fast-reroute
Displays Fast-reroute information.

igp-shortcut
Displays IS-IS IGP shortcut routing table information.

prefix
Displays IPv4 prefix information.

tilfa
Displays IPv4 prefix information.

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3. Added parameter algorithm in OcNOS version 6.6.0.

Example

```
ocnos#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 10:  VRF : default
Total number of routes: 3

 Destination         Metric      Next-Hop            Interface      Tag
C    2.2.2.2/32          10          --                  lo             0
L1   5.5.5.5/32          20          25.25.25.2          eth2           0
C    25.25.25.0/24       10          --                  eth2           0
```

Each entry in this table has a code preceding it, indicating the source of the routing entry.

The following table explains the fields in the output.

Table 53. 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).
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

The following table shows the route entry fields.

Table 54. route entry fields

Field	Description
Code	As explained in Table 53.
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.

© 2026 IP Infusion Inc. Proprietary

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

Execution modeand Privileged execution mode

Applicability

This command was introduced before OcNOS version 3.0.

Example

```
#show ip isis route fast-reroute
Tag   : 1  VRF : default
Total number of routes: 6

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

L2  1.1.1.1/32
    Primary Path via    : 172.27.1.2, xe15
    FRR Backup Path via : 172.26.1.1, xe20
    FRR Metric          : 50
    Protection Provided : LP NP BP DP

L2  3.3.3.3/32
    Primary Path via    : 172.27.1.2, xe15
    FRR Backup Path via : 172.26.1.1, xe20
    FRR Metric          : 60
    Protection Provided : LP BP

L2  4.4.4.4/32
    Primary Path via    : 172.28.1.1, xe3
    FRR Backup Path via : 172.27.1.2, xe15
    FRR Metric          : 30
    Protection Provided : LP BP

L2  7.7.7.77/32
    Primary Path via    : 172.28.1.1, xe3
    FRR Backup Path via : 172.27.1.2, xe15
    FRR Metric          : 30
    Protection Provided : LP BP
```

© 2026 IP Infusion Inc. Proprietary

```
L2  172.25.1.0/24
    Primary Path via    : 172.27.1.2, xe15
    FRR Backup Path via : 172.26.1.1, xe20
    FRR Metric          : 50
    Protection Provided : LP NP BP DP

L2  172.29.1.0/24
    Primary Path via    : 172.27.1.2, xe15
    FRR Backup Path via : 172.26.1.1, xe20
    FRR Metric          : 60
    Protection Provided : LP BP
```

© 2026 IP Infusion Inc. Proprietary

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

Execution modeand Privileged execution 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
```

Each entry in this table has a code preceding it, indicating the source of the routing entry. The following table shows these codes.

Table 55. 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).
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

© 2026 IP Infusion Inc. Proprietary

Table 55. Route codes and modifiers (continued)

Code	Description
	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 56. Route entry fields

Field	Description
Code	As explained in Table 53.
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.

© 2026 IP Infusion Inc. Proprietary

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 execution modeand Execution 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)
```

The following explains the fields shows for each route.

Table 57. 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.

© 2026 IP Infusion Inc. Proprietary

Table 57. show ip protocols output details (continued)

Field	Description
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.

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution 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 L1   40.40.40.0/24 [115/10] via 10.10.10.142, eth1, 00:00:50
                         [FRR-NH] via 30.30.30.144, eth3

i L1   50.50.50.0/24 [115/15] via 20.20.20.143, eth2, 00:00:50
                         [FRR-NH] via 10.10.10.142, eth1

i L1   60.60.60.0/24 [115/15] via 10.10.10.142, eth1, 00:00:50
                         [FRR-NH] via 20.20.20.143, eth2
```

© 2026 IP Infusion Inc. Proprietary

show ipv6 isis route

Use this command to display IPv6 IS-IS route information, including reachable prefixes, next-hop addresses, metrics, interfaces, and optional fast reroute or Flex-Algo data.

Command Syntax

```
show ipv6 isis route algorithm (128-255)(fast-reroute)
show ipv6 isis WORD topology (level-1|level-2|)
```

Parameters

128–255
Filters the route table to display only IPv6 IS-IS routes associated with a specific flex-Algo.

fast-reroute
Displays routes that have an associated fast reroute (FRR) backup path configured via mechanisms like TI-LFA.

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#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 10:  VRF : default
Total number of routes: 1

C    1111::1/128 [10]
      via ::, lo
```

The following table explains the output fields.

Table 58. show ipv6 isis topology output

Field	Description
Tag	Name that identifies the IS-IS area.
VRF	VRF name.
Prefix	IPv6 network prefix learned through IS-IS or configured locally (for example: connected).
Metric	ISIS metric used for SPF calculation (1-63).
Next-Hop	This route is available through the next hop router located at this IP address.
Interface	Interface from which the system was learned.

© 2026 IP Infusion Inc. Proprietary

Table 58. show ipv6 isis topology output (continued)

Field	Description
Level	IS-IS level from which the route was learned: L1 (Level-1) or L2 (Level-2).
Algo	Flexible Algorithm ID used for route calculation, if applicable.

© 2026 IP Infusion Inc. Proprietary

show ipv6 isis topology

Use this command to display IPv6 paths to Intermediate Systems.

Command Syntax

```
show ipv6 isis topology (level-1|level-2|)
show ipv6 isis WORD topology (level-1|level-2|)
```

Parameters

WORD
Display information for specified instance.

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

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ipv6 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
```

The following table explains the output fields.

Table 59. show ipv6 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).

© 2026 IP Infusion Inc. Proprietary

Table 59. show ipv6 isis topology output (continued)

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.

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution 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
```

© 2026 IP Infusion Inc. Proprietary

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 (level-1|level-2)
show isis database (detail|verbose) (level-1|level-2)
show isis database (detail|verbose) (level-1|level-2) WORD
show isis database WORD
show isis database WORD (level-1|level-2)
show isis database WORD (level-1|level-2) (detail|verbose)
show isis database WORD (detail|verbose)
show isis database WORD (detail|verbose) (level-1|level-2)
show isis database (level-1|level-2)
show isis database (level-1|level-2) (detail|verbose)
show isis database (level-1|level-2) (detail|verbose) WORD
show isis database (level-1|level-2) WORD
show isis database (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.

level-1
IS-IS level-1.

level-2
IS-IS level-2.

Command Mode

Execution mode and Privileged execution 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
```

The following table explains the output fields.

Table 60. 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.
P	Partition repair. A Level-1 area can become partitioned; this bit means the partition can be

Table 60. show isis database output (continued)

Field	Description
	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 61 if present in an LSP. For more about the TLV information, search for “IS-IS TLV Codepoints” on the Internet, check ISO/IEC 10589:2002(E), or other standard mentioned in the following table.

Table 61. 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

© 2026 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

Execution modeand Privileged execution 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
```

© 2026 IP Infusion Inc. Proprietary

```
#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
LDP-ISIS Sync Configured
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
```

The following table explains the output fields.

Table 62. 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.
LDP-ISIS Sync Configured	LDP IS-IS synchronization is enabled.
Holddown timer	Delay for notifications of LDP convergence to IS-IS
Remaining time	Remaining LDP convergence hold time in seconds.
Holddown timer not configured	The LDP convergence holddown timer has not been set.
Level-1 Metric	Interface metric value; used for SPF calculation.

© 2026 IP Infusion Inc. Proprietary

Table 62. show isis interface (continued)

Field	Description
Priority	Priority for designated IS election.
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

© 2026 IP Infusion Inc. Proprietary

show isis microloop-avoidance

Use this command to display microloop avoidance FSM states, timer values, and the oFIB route table.

Command Syntax

```
show isis microloop-avoidance ( algorithm|level-1|level-2|segment-routing) ( previous-info|detail|)
```

Parameters

algorithm
Display flexible algorithm

level-1
Display level-1 information only

level-2
Display level-2 information only

segment-routing
Display details of segment-routing

detail
Display detailed information for one or both levels

previous-info
Display details of previous oFIB execution

Command Mode

Privileged execution mode

Applicability

This command was introduced in OcNOS version 6.3.0. Added parameter algorithm in OcNOS version 6.6.0

Example

```
PE1#sh isis microloop-avoidance
Tag 1:  VRF : default
Level-1 status:
 FSM State: OFIB_STABLE

Level-2 status:
 FSM State: OFIB_STABLE

PE1#show isis microloop-avoidance detail
Tag 1:  VRF : default
Level-1 status:
 FSM State: OFIB_HOLDING_DOWN
 Event type: Neighbor Down
 Near end: 0000.0000.0001.00    Far end:0000.0000.0004.00
 Hold-down timer running: Yes    Time Remaining: 00:00:02.716
 Delay timer running: No

Level-2 status:
 FSM State: OFIB_STABLE
```

© 2026 IP Infusion Inc. Proprietary

```
oFIB Route Table:

     Destination         Metric      Next-Hop            Interface      Tag
L1   3.3.3.3/32          30          10.0.0.2            xe2            0
     Src: 0000.0000.0003
L1   4.4.4.4/32          40          10.0.0.2            xe2            0
     Src: 0000.0000.0004
L1   30.0.0.0/24         30          10.0.0.2            xe2            0
     Src: 0000.0000.0003
```

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)( algorithm)
```

Parameters

WORD
Routing area tag.

level-1
Displays Level 1 information only.

level-2
Displays Level 2 information only.

level-1-2
Displays both Level 1-2 information.

algorithm
Level 1-2

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 3.0. Added parameter algorithm in OcNOS version 6.6.0

Example

```
#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
```

© 2026 IP Infusion Inc. Proprietary

show isis topology

Use this command to display paths to Intermediate Systems.

Command Syntax

```
show isis topology ( algorithm|level-1|level-2|)
show isis WORD topology (algorithm|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

Execution modeand Privileged execution 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
```

The following table explains the output fields.

Table 63. show isis topology output

Field	Description
Tag	Name that identifies the IS-IS area.

Table 63. show isis topology output (continued)

Field	Description
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).
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.

© 2026 IP Infusion Inc. Proprietary

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

Execution 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
!
```

© 2026 IP Infusion Inc. Proprietary

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

Execution 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
!
```

© 2026 IP Infusion Inc. Proprietary

IS-IS Microloop-Avoidance Commands

This section describes the IS-IS microloop-avoidance commands:

microloop-avoidance 2070
microloop-avoidance flex-algo 2072
microloop-avoidance hold-timer 2073
microloop-avoidance max-fib 2074
microloop-avoidance segment-routing 2075
show isis microloop-avoidance 2076

© 2026 IP Infusion Inc. Proprietary

microloop-avoidance

Use this command to enable microloop avoidance using ordered FIB (oFIB) for an ISIS process.

Use the `no` form of this command to disable microloop avoidance for an ISIS process.

Microloop avoidance will be enabled with default hold-timer and max-fib values if it is not configured by the user:

- Default value of hold-timer is 6 seconds
- Default value for max-fib is 2 seconds

Microloop avoidance has to be enabled in conjunction with a fast reroute (FRR) mechanism such as TILFA that converts a sudden link or node failure into a non-urgent topology change by having complete repair path for all affected destinations. With microloop-avoidance, the convergence will be slower compared to traditional (SPF) convergence mechanisms. So, when there is any failure of a link or a set of links that are not protected using an FRR mechanism, it is advised to use conventional SPF (faster) mode of operation to minimize packet loss during convergence/reconvergence.

Command Syntax

```
microloop-avoidance(flex-algo|hold-timer|level-1|level-2|max-fib|segment-routing)
no microloop-avoidance(flex-algo|hold-timer|level-1|level-2|max-fib|segment-routing)
```

Parameters

flex-algo
Specifies the algorithm number used for Flexible Algorithm

hold-timer
Specifies the hold-timer value in seconds.

level-1
Applies the configuration to IS-IS Level-1 only.

level-2
Applies the configuration to IS-IS Level-2 only.

max-fib
Specifies the max-fib value for microloop-avoidance.

segment-routing
Enables Segment Routing functionality.

Default

None

Command Mode

ISIS router mode

Applicability

This command was introduced in OcNOS version 6.3.0. Added parameters flex-algo, hold-timer, max-fib, and segment-routing in OcNOS version 6.6.1.

© 2026 IP Infusion Inc. Proprietary

Examples

```
#configure terminal
(config)#router isis 100
(config-router)#microloop-avoidance level-1
(config-router)#microloop-avoidance level-2
(config-router)#

#configure terminal
(config)#router isis 100
(config-router)#no microloop-avoidance level-1
(config-router)#no microloop-avoidance level-2
(config-router)#
```

© 2026 IP Infusion Inc. Proprietary

microloop-avoidance flex-algo

Use this command to enable microloop avoidance using ordered FIB (oFIB) for an IS-IS Flexible Algorithm (Flex-Algo) instance.

Use the `no` form of this command to disable microloop avoidance for a Flex-Algo instance.

Microloop avoidance ensures loop-free convergence during topology changes by delaying FIB updates in a controlled sequence. When used with Flex-Algo, this feature applies the same loop-prevention behavior within the topology defined by that algorithm (for example: specific metric types, link affinities).

Command Syntax

```
microloop-avoidance flex-algo (128-255)(level-1|level-2)
no microloop-avoidance flex-algo (128-255)(level-1|level-2)
```

Parameters

128-255
Specifies the flex- algo instance range for which microloop-avoidance is to be enabled.

level-1
Applies microloop-avoidance to IS-IS Level-1 routing under the specified Flex-Algo.

level-2
Applies microloop-avoidance to IS-IS Level-2 routing under the specified Flex-Algo.

Default

None

Command Mode

ISIS router mode

Applicability

This command was introduced in OcNOS version 6.6.1.

Examples

```
#configure terminal
(config)#router isis 1
(config-router)#microloop-avoidance flex-algo 129 level-1
(config-router)#microloop-avoidance flex-algo 130 level-2

#configure terminal
(config)#router isis 1
(config-router)#no microloop-avoidance flex-algo 129 level-1
(config-router)#no microloop-avoidance flex-algo 130 level-2
```

© 2026 IP Infusion Inc. Proprietary

microloop-avoidance hold-timer

Use this command to configure the hold-timer value for microloop avoidance.

Use the `no` form of this command to return to the default hold-timer value for microloop avoidance.

Hold-timer value means the time that router will wait after receiving first LSP update with a topology change, to group all related LSP updates together to process using oFIB.

Microloop avoidance has to be enabled first for the particular level before configuring the hold-timer.

Hold-timer value must be configured same across all routers in the network to get the desired results.

Command Syntax

```
microloop-avoidance hold-timer (1-60) (level-1|level-2)
no microloop-avoidance hold-timer (level-1|level-2)
```

Parameters

1-60
Specifies the hold-timer value in seconds.

level-1
Applies the configuration to IS-IS Level-1 only.

level-2
Applies the configuration to IS-IS Level-1 only.

Default

Default value for hold-timer is 6 seconds

Command Mode

ISIS router mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
#configure terminal
(config)#router isis 100
(config-router)#microloop-avoidance hold-timer 3 level-1
(config-router)#microloop-avoidance hold-timer 3 level-2
(config-router)#

#configure terminal
(config)#router isis 100
(config-router)#no microloop-avoidance hold-timer level-1
(config-router)#no microloop-avoidance hold-timer level-2
(config-router)#
```

© 2026 IP Infusion Inc. Proprietary

microloop-avoidance max-fib

Use this command to configure the max-fib value for microloop avoidance.

Use the `no` form of this command to return to the default max-fib value for microloop avoidance.

The max-fib value is a network-wide constant that reflects the maximum time required to update a FIB irrespective of the change required. This value will be further multiplied by the oFIB RANK to get the actual FIB update delay time.

Microloop avoidance has to be enabled first for the particular level before configuring max-fib.

Max-fib value must be configured same across all routers in the network to get the desired results.

Command Syntax

```
microloop-avoidance max-fib (1-60) (level-1|level-2)
no microloop-avoidance max-fib (level-1|level-2)
```

Parameters

1-60
Specifies the max-fib value in seconds

level-1
Applies the configuration to IS-IS Level-1 only.

level-2
Applies the configuration to IS-IS Level-2 only.

Default

2 seconds

Command Mode

ISIS router mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
#configure terminal
(config)#router isis 100
(config-router)#microloop-avoidance max-fib 3 level-1
(config-router)#microloop-avoidance max-fib 3 level-2
(config-router)#

#configure terminal
(config)#router isis 100
(config-router)#no microloop-avoidance max-fib level-1
(config-router)#no microloop-avoidance max-fib level-2
```

© 2026 IP Infusion Inc. Proprietary

microloop-avoidance segment-routing

Use this command to configure a SR micro-loop avoidance under IS-IS Segment Routing (SR) mode.

Use the `no` form of this command to disable microloop avoidance for under IS-IS Segment Routing (SR) mode.

When microloop avoidance is enabled under Segment Routing, the router ensures loop-free convergence by controlling FIB updates across the network during topology changes.

Command Syntax

```
microloop-avoidance segment-routing (level-1 | level-2)(proto ipv4|ipv6 fib-delay) <100-300000>
no microloop-avoidance segment-routing (level-1 | level-2)(proto ipv4|ipv6)
```

Parameters

level-1
Applies Segment Routing microloop-avoidance logic only to IS-IS Level-1.

level-2
Applies Segment Routing microloop-avoidance logic only to IS-IS Level-2.

prot0
Specifies the protocol family (IPv4 or IPv6) to which microloop avoidance will be applied.

ipv4
Specifies the IPv4 protocol family to which microloop avoidance will be applied.

ipv6
Specifies the IPv6 protocol family to which microloop avoidance will be applied.

100–300000
Sets the delay in milliseconds for updating the Forwarding Information Base (FIB). This delay helps prevent microloops by controlling the rate at which FIB changes are applied.

Default

If fib-delay is not configured, default fib-delay of 1500ms(1.5 sec) is configured.

Command Mode

ISIS router mode

Applicability

This command was introduced in OcNOS version 6.6.1.

Examples

```
#configure terminal
(config)#router isis 1
(config-router)#microloop-avoidance segment-routing level-1 proto ipv4 fib-delay 100
(config-router)#microloop-avoidance segment-routing level-2 proto ipv6 fib-delay 1000

#configure terminal
(config)#router isis 1
(config-router)#no microloop-avoidance segment-routing level-1 proto ipv4
(config-router)#no microloop-avoidance segment-routing level-2 proto ipv6
```

© 2026 IP Infusion Inc. Proprietary

show isis microloop-avoidance

Use this command to display microloop avoidance FSM states, timer values, and the oFIB route table.

Command Syntax

```
show isis microloop-avoidance (flex-algo|level-1|level-2|segment-routing) (previous-info|detail|)
```

Parameters

flex-algo
Specifies the algorithm number used for Flexible Algorithm

detail
Display detailed information for one or both levels.

level-1
Filters the output to show microloop avoidance data for IS-IS Level-1 topology.

level-2
Filters the output to show microloop avoidance data for IS-IS Level-2 topology.

previous-info
Provides the details of previous oFIB executions.

segment-routing
Displays microloop avoidance configuration and state for IS-IS SR instances.

Command Mode

Privileged execution mode

Applicability

This command was introduced in OcNOS version 6.3.0. Added parameters flex-algo and previous-info in OcNOS version 6.6.1.

Example

```
ocnos#show isis microloop-avoidance detail
Tag 1:  VRF : default
Level-1 status:
 FSM State: OFIB_STABLE

Level-2 status:
 FSM State: OFIB_HOLDING_DOWN
 Event type: Neighbor Down
 Configured hold-timer: 6    Configured max-fib: 2
 Rank: 2
 Near end: 0000.0000.0004.00    Far end:0000.0000.0005.01
 Hold-down timer running: Yes    Time Remaning: 00:00:03.644
 Delay timer running: No

oFIB Route Table:

 Destination         Metric      Next-Hop            Interface      Tag

L2   2.2.2.2/32          30          3.5.0.3             eth2            0
```

```
     Src: 0000.0000.0002
L2   2.4.0.0/24          30          3.5.0.3             eth2            0
     Src: 0000.0000.0002
L2   2.6.0.0/24          30          3.5.0.3             eth2            0
     Src: 0000.0000.0002
L2   3.4.0.0/24          40          3.5.0.3             eth2            0
     Src: 0000.0000.0004
L2   4.4.4.4/32          40          3.5.0.3             eth2            0
     Src: 0000.0000.0004

ocnos# show isis microloop-avoidance previous-info
Tag 1:  VRF : default
Level-2 status:
 Last Event type: Neighbor Down
 Started at: 25-Jun-2025 13:23:51.873,  Completed at: 25-Jun-2025 13:23:57.873
 Near end: 0000.0000.0044.02,  Far end: 0000.0000.0044.02
 Hold-timer: 6s,  Max-fib: 2s
 Computed Rank: 0,  Delay time: 00:00:06

Total Delayed Routes Count: 2
Last Delayed Routes:
 4.4.4.4/32
 7.7.7.77/32

ocnos#show isis microloop-avoidance previous-info
Tag 1:  VRF : default
Level-2 status:
 Last Event type: Neighbor Up
 Started at: 25-Jun-2025 13:24:16.693,  Completed at: 25-Jun-2025 13:24:23.694
 Near end: 0000.0000.0044.02,  Far end: 0000.0000.0044.02
 Hold-timer: 6s,  Max-fib: 2s
 Computed Rank: 1,  Delay time: 00:00:08

Total Delayed Routes Count: 2
Last Delayed Routes:
 4.4.4.4/32
 7.7.7.77/32

ocnos#show isis microloop-avoidance segment-routing info
IS-IS level-1 proto-IPv4, microloop-avoidance segment-routing is enabled
  fib-delay 10000 ms.
  algorithm_num: 0
microloop-avoidance segment-routing SPF stage: Active
 Route:44.44.44.44/32
 Eligible vertex:PE2.00-00
        pnode: 0000.0001.0030.00
        qnode: 0000.0001.0044.00
        sid list:
                sid: 16030
                sid: 24320
        next-hop:
                ifindex: 10101
                ip_addr: 10.1.1.2
```

LAYER 3 SUB-INTERFACE CONFIGURATION

Layer 3 Sub-interface Configuration 2079
Topology 2079
Creating a Sub-interface 2079

© 2026 IP Infusion Inc. Proprietary

Layer 3 Sub-interface Configuration

This section contains examples of configuring sub-interfaces.

A single physical interface when required to handle multiple VLAN traffic, can be divided into multiple logical interfaces called sub-interfaces.

All sub-interfaces under a physical port will use their parent port for transmitting and receiving data.

Sub-interfaces can be used for various purposes, as for inter-vlan routing to happen when router has only one physical interface, two sub-interfaces each with different IP network can be created under it and data can be routed between them.

Sub-interfaces let you divide a physical interface into multiple logical interfaces that are tagged with different VLAN identifiers. Because VLANs allow you to keep traffic separate on a given physical interface, you can increase the number of interfaces available to your network without adding additional physical interfaces.

Note: Refer to the release note for features supported by L3 Sub-interface.

Topology

Figure 141 shows and example of sub-interface configuration. In this example, there are two routers, R1 and R2, and the eth1 interface of R1 is connected directly to eth2 of R2 using an Ethernet cable.

Figure 141. Subinterface connections

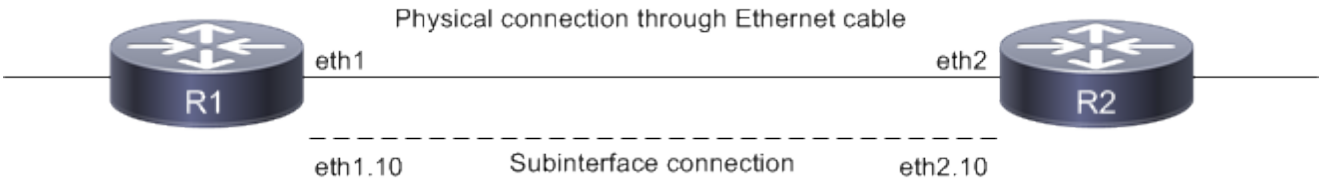

The eth1.10 subinterface is created on R1, and eth2.10 is created on R2.

Note: Layer 3 Subinterfaces can be created on physical and LAG interfaces.

Creating a Sub-interface

`#configure terminal`	Enter configure mode.
`(config)#interface eth1`	Enter interface mode
`(config-if)#interface eth1.10`	Creates a sub-interface as eth1.10
`(config-if)#encapsulation dot1q 10`	Configure the encapsulation as dot1q matching VLAN 10
`(config-if)#ip address 10.10.10.1/24`	Assigning IP address to sub-interface
`(config-if)#commit`	Commit the candidate configuration to the running

© 2026 IP Infusion Inc. Proprietary

	configuration
(config-if)#exit	Exit interface mode

Creating a Sub-interface with Encapsulation

Double encapsulation as dot1q

#configure terminal	Configure terminal
(config)#interface eth1.1010	Configure sub-interface
(config-if)# encapsulation dot1q 10 inner-dot1q 10	Configure encapsulation with inner tag
(config-if)#ip address 192.168.1.50/24	Configuring ip address
(config-if)#commit	Commit the candidate configuration to the running configuration
(config-if)#exit	Exit configure mode

Double encapsulation as dot1ad

#configure terminal	Configure terminal
(config)#interface eth1.20	Configure sub-interface
(config-if)# encapsulation dot1ad 20 inner-dot1q 20	Configure encapsulation with inner tag
(config-if)#ip address 192.168.2.50/24	Configuring ip address
(config-if)#commit	Commit the candidate configuration to the running configuration
(config-if)#exit	Exit configure mode

Notes:

- Use dot1ad ethertype (0x8100 | 0x88a8 | 0x9100 | 0x9200) command to configure the service-tpid value on parent port of a sub-interface. By this the tpid used for service tag for a sub-interface may be inherited from the one applied to parent interface.
- For any dot1ad sub-interface to be functional, dot1ad ethertype should be set to desired value as 0x88a8/0x9100/0x9200. Default value is 0x8100. To verify the ethertype value for the interface use show interface <subinterface> command.

L3SI Statistics

Enable below commands to get sub-interface statistics

#configure terminal	Configure terminal
(config)#hardware-profile statistics ac-lif enable	Enable hardware profile statistics
(config)#commit	Commit the candidate configuration to the running configuration
(config)#exit	Exit configure mode

© 2026 IP Infusion Inc. Proprietary

Note: Reload the node, and then only statistics command will get effective.

Displaying Sub-interfaces

In OcNOS, sub-interfaces appear as any physical interface in the `show running-config` or the show ip interface brief output and can be configured as any other interface.

The following examples display subinterface information from various `show` commands.

Note: The below command output is just for reference and is not directly related to the configuration provided above.

show interface brief

```
RTR1#show interface brief

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,       PD(Min-links) - Protocol Down
Min-links
       DV - DDM Violation, NA - Not Applicable
       NOM - No operational members, PVID - Port Vlan-id
       HD  - ESI Hold Timer Down


--------------------------------------------------------------------------------
Ethernet   Type  PVID  Mode              Status  Reason  Speed Port
Interface                                                    Ch #
--------------------------------------------------------------------------------
ce49       ETH   --    routed            up      none    100g  --


--------------------------------------------------------------------------------
Interface     Type             Status  Reason  Speed
--------------------------------------------------------------------------------
ce49.2        SUBINTERFACE     up      --      0
ce49.3        SUBINTERFACE     up      --      0
ce49.4        SUBINTERFACE     up      --      0
ce49.5        SUBINTERFACE     up      --      0
ce49.6        SUBINTERFACE     up      --      0
```

show ip interface brief

```
RTR1#show ip interface brief

'*' - address is assigned by dhcp client

Interface            IP-Address      Admin-Status          Link-Status
ce49                 unassigned      up                    up
ce49.2               49.49.2.1       up                    up
ce49.3               49.49.3.1       up                    up
ce49.4               49.49.4.1       up                    up
ce49.5               49.49.5.1       up                    up
ce49.6               49.49.6.1       up                    up
```

show ip ospf neighbor with VRF enabled

```
RTR1#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(default):
```

© 2026 IP Infusion Inc. Proprietary

```
Neighbor ID     Pri   State           Dead Time   Address         Interface           Instance ID
4.4.4.4           1   Full/DR         00:00:32    48.48.2.2       vlan1.2                 0
4.4.4.4           1   Full/DR         00:00:38    48.48.3.2       vlan1.3                 0

Total number of full neighbors: 1
OSPF process 2 VRF(CUST-2):
Neighbor ID     Pri   State           Dead Time   Address         Interface           Instance ID
11.11.2.1         1   Full/DR         00:00:39    49.49.2.2       ce49.2                  0

Total number of full neighbors: 1
OSPF process 3 VRF(CUST-3):
Neighbor ID     Pri   State           Dead Time   Address         Interface           Instance ID
11.11.3.1         1   Full/Backup     00:00:33    49.49.3.2       ce49.3                  0

Total number of full neighbors: 1
OSPF process 4 VRF(CUST-4):
Neighbor ID     Pri   State           Dead Time   Address         Interface           Instance ID
11.11.4.1         1   Full/Backup     00:00:31    49.49.4.2       ce49.4                  0

Total number of full neighbors: 1
OSPF process 5 VRF(CUST-5):
Neighbor ID     Pri   State           Dead Time   Address         Interface           Instance ID
11.11.5.1         1   Full/Backup     00:00:39    49.49.5.2       ce49.5                  0
```

show ip route with VRF enabled

```
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            1.2.200.0/24 is directly connected, xe1.200, 01:29:19
O            4.4.4.4/32 [110/11] via 48.48.3.2, vlan1.3, 00:37:17
                        [110/11] via 48.48.2.2, vlan1.2
O            44.44.44.0/24 [110/2] via 48.48.3.2, vlan1.3, 00:37:17
                           [110/2] via 48.48.2.2, vlan1.2
C            47.47.2.0/24 is directly connected, xe47.2, 00:34:42
C            48.48.2.0/24 is directly connected, vlan1.2, 00:41:19
C            48.48.3.0/24 is directly connected, vlan1.3, 00:41:19
C            127.0.0.0/8 is directly connected, lo, 01:30:09

Gateway of last resort is not set

RTR1#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            1.2.200.0/24 is directly connected, xe1.200, 01:29:32
O            4.4.4.4/32 [110/11] via 48.48.3.2, vlan1.3, 00:37:30
                        [110/11] via 48.48.2.2, vlan1.2
O            44.44.44.0/24 [110/2] via 48.48.3.2, vlan1.3, 00:37:30
                           [110/2] via 48.48.2.2, vlan1.2
C            47.47.2.0/24 is directly connected, xe47.2, 00:34:55
C            48.48.2.0/24 is directly connected, vlan1.2, 00:41:32
```

```
C            48.48.3.0/24 is directly connected, vlan1.3, 00:41:32
C            127.0.0.0/8 is directly connected, lo, 01:30:22
IP Route Table for VRF "management"
C            127.0.0.0/8 is directly connected, lo.management, 01:30:22
C            192.168.10.0/24 is directly connected, eth0, 01:30:22
IP Route Table for VRF "CUST-1"
C            127.0.0.0/8 is directly connected, lo.CUST-1, 01:30:22
IP Route Table for VRF "CUST-2"
C            1.1.2.0/24 is directly connected, xe1.2, 01:29:35
C            1.2.101.0/24 is directly connected, xe1.101, 01:29:34
C            1.3.201.0/24 is directly connected, xe1.201, 01:29:32
O            11.11.2.0/24 [110/20] via 49.49.2.2, ce49.2, 00:51:06
O            11.12.101.0/24 [110/20] via 49.49.2.2, ce49.2, 00:51:06
O            11.13.201.0/24 [110/20] via 49.49.2.2, ce49.2, 00:51:06
C            49.49.2.0/24 is directly connected, ce49.2, 01:29:31
C            127.0.0.0/8 is directly connected, lo.CUST-2, 01:30:22
IP Route Table for VRF "CUST-3"
C            1.1.3.0/24 is directly connected, xe1.3, 01:29:35
C            1.2.102.0/24 is directly connected, xe1.102, 01:29:34
C            1.3.202.0/24 is directly connected, xe1.202, 01:29:32
O            11.11.3.0/24 [110/20] via 49.49.3.2, ce49.3, 01:12:44
O            11.12.102.0/24 [110/20] via 49.49.3.2, ce49.3, 01:12:44
O            11.13.202.0/24 [110/20] via 49.49.3.2, ce49.3, 01:12:44
C            49.49.3.0/24 is directly connected, ce49.3, 01:29:31
```

© 2026 IP Infusion Inc. Proprietary

LAYER 3 SUB-INTERFACE COMMANDS

Layer 3 Sub-interface Commands 2085
encapsulation 2086
interface IFNAME.SUBINTERFACE_ID 2087
show interface IFNAME.SUBINTERFACE_ID 2088

© 2026 IP Infusion Inc. Proprietary

Layer 3 Sub-interface Commands

This section describes the Layer 3 sub-interface commands:

encapsulation 2086
interface IFNAME.SUBINTERFACE_ID 2087
show interface IFNAME.SUBINTERFACE_ID 2088

© 2026 IP Infusion Inc. Proprietary

encapsulation

Use this command to configure encapsulation under a sub-interface. Using this command, a Layer 3 sub-interface can be configured as a port-vlan or stacked vlan. Before configuring the encapsulation on sub-interface, the operating state of the sub-interface is admin down. After configuring the encapsulation, the operating state of the sub-interface becomes up.

Note: Due to hardware solution restriction,MPLS termination on double-tagged sub-interface is not supported for Qumran1 (Q1) series platforms devices.

Command Syntax

```
encapsulation ((dot1q|dot1ad) VLAN_ID (inner-dot1q VLAN_ID|))
no encapsulation
```

Parameters

dot1q
IEEE 802.1Q VLAN-tagged packets

dot1ad
IEEE 802.1ad VLAN-tagged packets

VLAN_ID
First (outer) VLAN identifier on the subinterface <2-4094>

inner-dot1q VLAN_ID
Second (inner 802.1Q) VLAN identifier on the subinterface <2-4094>

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS version 3.0. For DC - OcNOS version 7.0.0 is the applicability.

Example

```
#configure terminal
(config-if)#interface xe9.1
(config-if)#encapsulation dot1q 10

(config-if)#interface xe9.2
(config-if)#encapsulation dot1q 20 inner-dot1q 20

(config-if)#interface xe9.3
(config-if)#encapsulation dot1ad 30

(config-if)#interface xe9.4
(config-if)#encapsulation dot1ad 40 inner-dot1q 40
```

© 2026 IP Infusion Inc. Proprietary

interface IFNAME.SUBINTERFACE_ID

Use this command to configure a sub-interface. Sub-interfaces are supported for raw Ethernet interfaces as well as dynamic/static LAG interfaces. The number of sub-interfaces supported is 2000 per parent port and a total of 4000 sub-interfaces per device. This command displays an error when the identifier of the sub-interface is out of range.

Use `no` form of this command to unconfigure a sub-interface.

Command Syntax

```
interface IFNAME
no interface IFNAME
```

Parameters

IFNAME

Specifies the sub-interface name, represented in the format interface-name.sub-interface ID (for example, xe1.5, where xe1 is the interface name and 5 is the sub-interface ID). The valid sub-interface range is <1-2000> .

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS version 3.0. For DC - OcNOS version 7.0.0 is the applicability.

Example

```
(config)#int xe1.5
(config-if)#exit
(config)#no interface xe1.5

(config)#int po1
(config-if)#exit
(config)#int po1.1
(config-if)#exit
(config)#no interface po1.1

(config)#int sa1
(config-if)#exit
(config)#int sa1.1
(config-if)#exit
(config)#no interface sa1.1
```

© 2026 IP Infusion Inc. Proprietary

show interface IFNAME.SUBINTERFACE_ID

Use this command to display the details of the sub-interface.This command displays the information about the operating status, hardware address, VRF binding details, and input/output counters. This command can display details of the a sub-interface for a dynamic/static LAG as well.

Command Syntax

```
show interface IFNAME
```

Parameters

IFNAME

Specifies the sub-interface name, represented in the format interface-name.sub-interface ID (for example, xe1.5, where xe1 is the interface name and 5 is the sub-interface ID). The valid sub-interface range is <1-2000> .

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 3.0. For DC - OcNOS version 7.0.0 is the applicability.

Example

```
#show int xe1.1
Interface xe1.1
  Hardware is SUBINTERFACE  Current HW addr: 6cb9.c500.1647
  Physical:6cb9.c500.1647  Logical:(not set)
  Port Mode is Router
  Interface index: 20482049
  Metric 1 mtu 1500
  <UP,BROADCAST,RUNNING,MULTICAST>
  VRF Binding: Not bound
  Label switching is disabled
  No Virtual Circuit configured
  Administrative Group(s): None
  DHCP client is disabled.
  Last Flapped: Never
  Statistics last cleared: Never
  inet6 fe80::6eb9:c5ff:fe00:1647/64
  RX
    unicast packets 0 multicast packets 0 broadcast packets 0
    input packets 0 bytes 0
    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 0 broadcast packets 0
    output packets 0 bytes 0
    jumbo packets 0
    output errors 0 collision 0 deferred 0 late collision 0
    output discard 0
```

© 2026 IP Infusion Inc. Proprietary

```
    Tx pause 0

#show int sa1.1
Interface sa1.1
  Hardware is SUBINTERFACE  Current HW addr: 6cb9.c500.1647
  Physical:6cb9.c500.1647  Logical:(not set)
  Port Mode is Router
  Interface index: 409602049
  Metric 1 mtu 1500
  <UP,BROADCAST,RUNNING,MULTICAST>
  VRF Binding: Not bound
  Label switching is disabled
  No Virtual Circuit configured
  Administrative Group(s): None
  DHCP client is disabled.
  Last Flapped: Never
  Statistics last cleared: Never
  inet6 fe80::6eb9:c5ff:fe00:1647/64
  RX
    unicast packets 0 multicast packets 0 broadcast packets 0
    input packets 0 bytes 0
    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 0 broadcast packets 0
    output packets 0 bytes 0
    jumbo packets 0
    output errors 0 collision 0 deferred 0 late collision 0
    output discard 0
    Tx pause 0

#show int po1.3
Interface po1.3
  Hardware is SUBINTERFACE  Current HW addr: 0030.abf1.0ec8
  Physical:0030.abf1.0ec8  Logical:(not set)
  Port Mode is Router
  Interface index: 204802051
  Metric 1 mtu 1500
  <UP,BROADCAST,RUNNING,MULTICAST>
  VRF Binding: Not bound
  Label switching is disabled
  No Virtual Circuit configured
  Administrative Group(s): None
  DHCP client is disabled.
  Last Flapped: Never
  Statistics last cleared: Never
  inet 23.0.0.2/24 broadcast 23.0.0.255
  inet6 fe80::230:abff:fef1:ec8/64
  RX
    unicast packets 0 multicast packets 0 broadcast packets 0
    input packets 141805 bytes 9643544
    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 0 broadcast packets 0
    output packets 0 bytes 0
    jumbo packets 0
    output errors 0 collision 0 deferred 0 late collision 0
    output discard 0
    Tx pause 0
```

NEIGHBOR DISCOVERY CONFIGURATION

Neighbor Discovery Configuration 2091
ARP/Neighbor Discovery Operation 2091

Neighbor Discovery Configuration

This section provides an overview of Neighbor Discovery (ND) configuration.

The Address Resolution Protocol (ARP) translates network layer addresses into link-layer addresses. ARP converts a an IPv4 address to an Ethernet address (MAC address).

In Internet Protocol Version 6 (IPv6) networks, the functionality of ARP is provided by the Neighbor Discovery (ND) protocol. Neighbor Discovery operates at the link layer and is responsible for auto configuration of nodes, discovery of other nodes on the link, determining the link layer addresses of other nodes, duplicate address detection, finding available routers and Domain Name System (DNS) servers, address prefix discovery, and maintaining reachability information about the paths to other active neighbor nodes.

ARP/Neighbor Discovery Operation

Neighbor Discovery module manages the ARP and IPv6 ND entries and provides information to other protocols, the forwarding module for their use, and reports state changes.

Neighbor Entry States

INCOMPLETE	Address resolution is in progress and the link-layer address of the neighbor has not yet been determined.
REACHABLE	The neighbor is known to have been reachable recently.
STALE	The neighbor is no longer known to be reachable (not used by higher level protocol for reachable-time) but still valid and used for forwarding. Until any control packet is sent to the neighbor, no attempt will be made to verify its reachability.
DELAY	The neighbor is no longer known to be reachable, and traffic has recently been sent to the neighbor. Rather than probing the neighbor immediately, delay sending probes for a short while to give upper-layer protocols a chance to provide reachability confirmation.
PROBE	The neighbor is no longer known to be reachable, and probes are being sent to verify reachability.
FAILED	

Below are timers that control the above state transitions.

reachable-time	The amount of time the entry is in REACHABLE state. The default value is 60 seconds. Once a neighbor is found, the entry is considered reachable for at least a random value between (A) and (3xA) where: A = reachable-time /2 Once entry reachability expires it moves to STALE state. The entry's reachability is extended if it receives positive feedback (ARP reply/NA).

© 2026 IP Infusion Inc. Proprietary

stale-time	Determines how often to check for stale neighbor entries. The default value is 1440 seconds. Once the stale-timer expires, the entry is marked for garbage collection.
arp-aging-time	The approximate amount of time an ARP entry is valid. The default value is 1500 seconds. Aging time is configured internally as: aging-time = reachable-time + stale-time

The garbage collector runs every 60 seconds once, to clean-up the entries which have crossed STALE timeout and FAILED entries. For improved performance, the following are the additional criteria for the neighbor entries to be garbage collected (removal of the entry completely).

- The garbage collector does not run if the total number of entries is less than 2048 that conveys some STALE, and FAILED entries still exist. The entry is refreshed ONLY when higher-level protocols use it.
- The garbage collector always runs if the number of entries is more than 262144 for IPv4 and 131072 for IPv6.

The garbage collection triggers the refresh of neighbor entries which are marked to be garbage collected. A total of 3 retries with a 3 second interval is done to reach the neighbor, before declaring the neighbor as not reachable by the forwarder. During this retry period, neighbor entries are still in use for hardware forwarding. If the neighbor is still not reachable after retries (after 9 seconds), the neighbor entry is removed from hardware forwarding as well.

Configuring ARP for IPv4

The procedures in this section use the topology in Figure 142

Figure 142. ARP for IPv4

RTR1

`#configure terminal`	Enter the configure mode.
`(config)#interface xe1`	Enter interface mode.
`(config-if)#ip address 2.2.2.2/24`	Configure IP address on the interface.
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode.
`(config)#arp 2.2.2.3 0000.0000.0003`	Configure ARP entry for neighbor.
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit configure mode.

© 2026 IP Infusion Inc. Proprietary

RTR2

`#configure terminal`	Enter the configure mode.
`(config)#interface xe1`	Enter interface mode.
`(config-if)#ip address 2.2.2.3/24`	Configure IP address on the interface.
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode.
`(config)#arp 2.2.2.2 0000.0000.0004`	Configure ARP entry for neighbor.
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit configure mode.

Validation

```
#show arp
Address                  HWaddress             Interface      Type
2.2.2.3                  00:00:00:00:00:03      eth1          Static
10.12.65.1               00:07:7d:67:6f:bf      eth0          Dynamic
```

Configuring Neighbor Discovery for IPv6

The procedures in this section use the topology in Figure 143.

Figure 143. ND for IPv6

RTR1

`#configure terminal`	Enter the configure mode.
`(config)#interface xe1`	Enter interface mode.
`(config-if)#ipv6 address 3ffe:506::1/48`	Configure IPv6 address on the interface.
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode.
`(config)# ipv6 neighbor 3ffe:506::2 xe1 0000.0000.0004`	Configure neighbor IPv6 address and MAC.
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit configure mode.

© 2026 IP Infusion Inc. Proprietary

RTR2

`#configure terminal`	Enter the configure mode.
`(config)#interface xe1`	Enter interface mode.
`(config-if)#ipv6 address 3ffe:506::2/48`	Configure IPv6 address on the interface.
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode.
`(config)# ipv6 neighbor 3ffe:506::1 xe1 0000.0000.0003`	Configure neighbor IPv6 address and MAC.
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config)#exit`	Exit configure mode.

Validation

```
#show ipv6 neighbors
IPv6 Address            MAC Address      Interface        Type
3ffe:506::2             0000.0000.0004    xe1             static
```

NEIGHBOR DISCOVERY COMMAND REFERENCE

This section describes the Neighbor Discovery commands:

Neighbor Discovery Command Reference 2096
- arp-ageing-timeout 2097
- arp-reachable-time 2098
- clear arp 2099
- clear ipv6 neighbors 2100
- debug ip arp 2101
- debug ipv6 nd 2102
- ip arp 2103
- ip arp vrf 2104
- ip proxy-arp 2105
- ipv6 nd current-hoplimit 2106
- ipv6 nd link-mtu 2107
- ipv6 nd managed-config-flag 2108
- ipv6 nd other-config-flag 2109
- ipv6 nd prefix 2110
- ipv6 nd ra-interval 2112
- ipv6 nd ra-lifetime 2113
- ipv6 nd reachable-time 2114
- ipv6 nd retransmission-time 2115
- ipv6 nd suppress-ra 2116
- ipv6 neighbor 2117
- nd-ageing-timeout 2118
- nd-reachable-time 2119
- no debug all 2120
- show arp 2121
- show debugging ip arp 2123
- show debugging ipv6 nd 2124
- show ipv6 neighbors 2125

© 2026 IP Infusion Inc. Proprietary

Neighbor Discovery Command Reference

This section describes the Neighbor Discovery commands:

arp-ageing-timeout 2097
arp-reachable-time 2098
clear arp 2099
clear ipv6 neighbors 2100
debug ip arp 2101
debug ipv6 nd 2102
ip arp 2103
ip arp vrf 2104
ip proxy-arp 2105
ipv6 nd current-hoplimit 2106
ipv6 nd link-mtu 2107
ipv6 nd managed-config-flag 2108
ipv6 nd other-config-flag 2109
ipv6 nd prefix 2110
ipv6 nd ra-interval 2112
ipv6 nd ra-lifetime 2113
ipv6 nd reachable-time 2114
ipv6 nd retransmission-time 2115
ipv6 nd suppress-ra 2116
ipv6 neighbor 2117
nd-ageing-timeout 2118
nd-reachable-time 2119
no debug all 2120
show arp 2121
show debugging ip arp 2123
show debugging ipv6 nd 2124
show ipv6 neighbors 2125

© 2026 IP Infusion Inc. Proprietary

arp-ageing-timeout

Use this command to set the ARP aging timeout.

Note: The bridge aging time affects the ARP entries which are dependent upon the MAC addresses in hardware. If a MAC address ages out, it causes the corresponding ARP entry to refresh.

Use `no` form of this command to set the aging to its default value.

Command Syntax

```
arp-ageing-timeout <60-36000>
no arp-ageing-timeout
```

Parameter

<60-36000>
ARP aging timeout.

Default value

By default, the ARP aging timeout is 1500.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface xe2
(config-if)#arp-ageing-timeout 5000
(config-if)#exit

(config)#interface xe2
(config-if)#no arp-ageing-timeout
```

© 2026 IP Infusion Inc. Proprietary

arp-reachable-time

Use this command to set the ARP reachable time.

Use `no` form of this command to set the reachable time to its default value.

Command Syntax

```
arp-reachable-time <10-36000>
no arp-reachable-time
```

Parameters

<10-36000>
Specify the ARP reachable time.

Default value

By default, arp-reachable-time is 60

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface xe2
(config-if)#arp-reachable-time 120
(config-if)#exit

(config)#interface xe2
(config-if)#no arp-reachable-time
```

© 2026 IP Infusion Inc. Proprietary

clear arp

Use this command to clear dynamic ARP entries.

Command Syntax

```
clear arp (| A.B.C.D) (| vrf (all | VRFNAME | default))
clear arp IFNAME (| vrf (all | VRFNAME | default))
```

Parameters

A.B.C.D
Specify the IP address of the ARP entry.

IFNAME
Specify the name of the interface.

all
All VRFs.

VRFNAME
VPN routing/forwarding instance name.

default
Default VRF.

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear arp
#clear arp xe1
#clear arp 10.10.10.10
#clear arp vrf VRF1
```

© 2026 IP Infusion Inc. Proprietary

clear ipv6 neighbors

Use this command to clear dynamic neighbor entries.

Command Syntax

```
clear ipv6 neighbors (|X:X::X:X) (|vrf (all | VRFNAME | default))
clear ipv6 neighbors IFNAME (|vrf (all | VRFNAME | default))
```

Parameters

X:X::X:X
Specify the neighbor's IPv6 address.

IFNAME
Specify the name of the interface.

all
All VRFs.

VRFNAME
VPN routing/forwarding instance name.

default
Default VRF.

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear ipv6 neighbors
#clear ipv6 neighbors xe1
#clear ipv6 neighbors 2000::3
#clear ipv6 neighbors vrf VRF1
```

© 2026 IP Infusion Inc. Proprietary

debug ip arp

Use this command to enable debugging for ARP events.

Note: Packet dumps are not available as part of debug logs. Use the `tcpdump` command to view packets.

Use the `no` parameter with this command to disable event debugging.

Command Syntax

```
debug ip arp event
no debug ip arp event
```

Parameters

event
ARP event debugging.

Command Mode

Exec mode and configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#debug ip arp event
```

debug ipv6 nd

Use this command to enable debugging for neighbor events.

Note: Packet dumps are not available as part of debug logs. Use the `tcpdump` command to view packets.

Use the `no` parameter with this command to disable event debugging.

Command Syntax

```
debug ipv6 nd event
no debug ipv6 nd event
```

Parameters

event
ND event debugging.

Command Mode

Exec mode and configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#debug ipv6 nd event
```

© 2026 IP Infusion Inc. Proprietary

ip arp

Use this command to create a static ARP entry.

Use the `no` parameter to remove the static ARP entry.

Command Syntax

```
ip arp A.B.C.D XXXX.XXXX.XXXX (alias|)
no ip arp A.B.C.D
```

Parameters

A.B.C.D
Specify the IP address of the ARP entry.

XXXX.XXXX.XXXX
Specify the MAC (hardware) address of the ARP entry.

alias
Specify the response to ARP requests for the IP address.

Default value

No default value is specified

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#ip arp 10.10.10.10 0000.0001.4566
(config)#no ip arp 10.10.10.10
```

© 2026 IP Infusion Inc. Proprietary

ip arp vrf

Use this command to create a static ARP entry for the non-default VRF.

Note: This command is supported only for multiple VRF support.

Command Syntax

```
ip arp vrf NAME A.B.C.D XXXX.XXXX.XXXX (alias|)
```

Parameter

NAME
Specify VRF name to which entry need to be added.

A.B.C.D
Specify the IP address of the ARP entry.

XXXX.XXXX.XXXX
Specify the MAC (hardware) address of the ARP entry.

alias
Specify the response to ARP requests for the IP address.

Default value

No default value is specified

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#ip arp vrf VRF1 10.10.10.10 0000.0001.4566
```

© 2026 IP Infusion Inc. Proprietary

ip proxy-arp

Use this command to enable the proxy ARP feature.

Use the `no` parameter to disable the proxy ARP feature.

Command Syntax

```
ip proxy-arp
no ip proxy-arp
```

Parameter

None

Default value

No default value is specified

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface xe1
(config-if)#ip proxy-arp
(config-if)#no ip proxy-arp
```

© 2026 IP Infusion Inc. Proprietary

ipv6 nd current-hoplimit

Use this command to set an ND (Neighbor Discovery) advertised hop limit for an interface.

Use the `no` option of this command to set the current hop limit to its default value.

Command Syntax

```
ipv6 nd current-hoplimit <0-255>
no ipv6 nd current-hoplimit
```

Parameter

<0-255>
Hop limit.

Default

By default, hop limit is 64

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth3
(config-if)#ipv6 nd current-hoplimit 10
(config-if)#no ipv6 nd current-hoplimit
```

© 2026 IP Infusion Inc. Proprietary

ipv6 nd link-mtu

Use this command to set an advertised maximum transmission unit (MTU).

Use the `no` option with the command to reset the MTU to the default statute.

Command Syntax

```
ipv6 nd link-mtu (<1280-65535>)
no ipv6 nd link-mtu
```

Parameters

<1280-65535>
Link MTU value.

Default

By default, link MTU value is 1500

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth3
(config-if)#ipv6 nd link-mtu 1600
(config-if)#no ipv6 nd link-mtu
```

© 2026 IP Infusion Inc. Proprietary

ipv6 nd managed-config-flag

Use this command to set the managed address configuration flag in the Router Advertisement to be used for the IPv6 address auto-configuration.

Use the `no` parameter with this command to reset the value to default.

Command Syntax

```
ipv6 nd managed-config-flag
no ipv6 nd managed-config-flag
```

Parameters

None

Default

The managed address configuration flag is not set.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth3
(config-if)#ipv6 nd managed-config-flag
(config-if)#no ipv6 nd managed-config-flag
```

© 2026 IP Infusion Inc. Proprietary

ipv6 nd other-config-flag

Use this command to set the other stateful configuration flag in Router Advertisement to be used for IPv6 address auto-configuration.

Use `no` parameter with this command to reset the value to default.

Command Syntax

```
ipv6 nd other-config-flag
no ipv6 nd other-config-flag
```

Parameters

None

Default

Other stateful configuration flag is not set.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth3
(config-if)#ipv6 nd other-config-flag
```

© 2026 IP Infusion Inc. Proprietary

ipv6 nd prefix

Use this command to set IPv6 prefix information that is advertised for address auto-configuration.

Use `no` parameter with this command to remove an IPv6 prefix.

Command Syntax

```
ipv6 nd prefix X:X::X:X/M <0-4294967295> <0-4294967295> (off-link|) (no-autoconfig|)
ipv6 nd prefix X:X::X:X/M
ipv6 nd prefix valid-lifetime <0-4294967295>
ipv6 nd prefix preferred-lifetime <0-4294967295>
ipv6 nd prefix offlink
ipv6 nd prefix no-autoconf
no ipv6 nd prefix X:X::X:X/M (off-link|) (no-autoconfig|)
no ipv6 nd prefix valid-lifetime
no ipv6 nd prefix preferred-lifetime
```

Parameters

`X:X::X:X/M`
IPv6 prefix.

`<0-4294967295>`
Valid lifetime in seconds. Use the preferred-lifetime value < valid-lifetime.

`<0-4294967295>`
Preferred lifetime in seconds. Use the preferred-lifetime value < valid-lifetime.

`off-link`
Do not use prefix for onlink determination.

`no-autoconfig`
Do not use prefix for autoconfiguration.

Default

By default, valid life time is 2592000 seconds and preferred life time is 604800 seconds

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Defaults

The default valid lifetime is 2592000 seconds.

The default preferred lifetime is 604800 seconds.

Examples

```
(config)#interface eth3
(config-if)#ipv6 nd prefix 2001:ffff::/64
```

© 2026 IP Infusion Inc. Proprietary

```
(config)#interface eth3
(config-if)#ipv6 nd prefix no-autoconf

(config)#interface eth3
(config-if)#ipv6 nd prefix preferred-lifetime 550000
```

© 2026 IP Infusion Inc. Proprietary

ipv6 nd ra-interval

Use this command to specify the interval between IPv6 Router Advertisements (RA).

Use `no` parameter with this command to set the value to its default.

Note: The RA interval will have a random variation from configured interval to avoid synchronization of advertisement with routers.

Command Syntax

```
ipv6 nd ra-interval <4-1800> (<3-1350>|)
no ipv6 nd ra-interval
```

Parameter

<4-1800>
RA interval in seconds.

<3-1350>
Minimum RA interval (in seconds).

Default

By default, RA interval is 600 seconds, and Minimum RA interval is 0.33*RA interval.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth3
(config-if)#ipv6 nd ra-interval 60 15
(config-if)#ipv6 nd prefix 3ffe:ffff:ffff::/64

(config-if)#no ipv6 nd ra-interval
```

© 2026 IP Infusion Inc. Proprietary

ipv6 nd ra-lifetime

Use this command to specify the Router Advertisement (RA) lifetime for this device to act as a default gateway for the network.

Use `no` parameter with this command to reset the value to default.

Command Syntax

```
ipv6 nd ra-lifetime <0-9000>
no ipv6 nd ra-lifetime
```

Parameter

<0-9000>
RA lifetime duration in seconds.

Default

The default RA lifetime is 1800 seconds.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth3
(config-if)#ipv6 nd ra-lifetime 9000
(config-if)#no ipv6 ra-lifetime 9000
```

© 2026 IP Infusion Inc. Proprietary

ipv6 nd reachable-time

Use this command to specify the reachable time in the Router Advertisement to be used for detecting unreachability of the IPv6 neighbor.

Use the `no` parameter with this command to set the value to its default.

Command Syntax

```
ipv6 nd reachable-time <0-3600000>
no ipv6 nd reachable-time
```

Parameter

<0-3600000>
Reachable time in milliseconds.

Default

By default, reachable time is zero (0) milliseconds.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth3
(config-if)#ipv6 nd reachable-time 1800000
(config-if)#no ipv6 nd reachable-time 1800000
```

© 2026 IP Infusion Inc. Proprietary

ipv6 nd retransmission-time

Use this command to set an IPv6 advertised retransmission time for the current interface.

Use the `no` form of the command to set the retransmission time to its default value.

Command Syntax

```
ipv6 nd retransmission-time (0-4294967295)
no ipv6 nd retransmission-time
```

Parameter

<0-4294967295>
Retransmission time in milliseconds

Default

By default, retransmission time is zero (0) milliseconds.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth3
(config-if)#ipv6 nd retransmission-time 1200
(config-if)#no ipv6 nd retransmission-time
```

© 2026 IP Infusion Inc. Proprietary

ipv6 nd suppress-ra

Use this command to suppress IPv6 Router Advertisement (RA) transmission for the current interface. Router Advertisement is used for IPv6 stateless auto-configuration.

Use the `no` parameter with this command to enable Router Advertisement transmission.

Command Syntax

```
ipv6 nd suppress-ra (mtu |)
no ipv6 nd suppress-ra (mtu |)
```

Parameters

mtu
Send maximum transmission unit (MTU) in Router Advertisement messages

Default

By default, `ipv6 nd suppress-ra` is disabled.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth3
(config-if)#ipv6 nd suppress-ra
```

© 2026 IP Infusion Inc. Proprietary

ipv6 neighbor

Use this command to add a static neighbor entry.

Use the `no` form of this command to remove a static neighbor entry.

Command Syntax

```
ipv6 neighbor X:X::X:X IFNAME XXXX.XXXX.XXXX
no ipv6 neighbor X:X::X:X IFNAME
```

Parameter

X:X::X:X
Specify the neighbor's IPv6 address.

IFNAME
Specify the name of the interface.

xxxx.xxxx.xxxx

Specify the MAC hardware address.

Default value

No default value is specified

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#ipv6 neighbor 2000::3 xe1 0000.0002.3dc1
(config)#no ipv6 neighbor 2000::3 xe1
```

© 2026 IP Infusion Inc. Proprietary

nd-ageing-timeout

Use this command to set the neighbor ageing timeout value.

Use `no` form of this command to set the ageing to its default value.

Command Syntax

```
nd-ageing-timeout <60-36000>
no nd-ageing-timeout
```

Parameters

<60-36000>
ND ageing timeout in seconds.

IFNAME
Specify the name of the interface.

XXXX.XXXX.XXXX
Specify the MAC hardware address.

Default value

By default, nd-ageing-timeout is 1500 seconds

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface xe1
(config-if)#nd-ageing-timeout 3600
(config-if)#no nd-ageing-timeout
```

© 2026 IP Infusion Inc. Proprietary

nd-reachable-time

Use this command to set neighbor reachable time value.

Use `no` form of this command to set the reachable time to its default value.

Command Syntax

```
nd-reachable-time <10-36000>
no nd-reachable-time
```

Parameter

<10-36000>
ND reachable time in seconds.

Default value

By default, nd-reachable-time is 60 seconds

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface xe1
(config-if)#nd-reachable-time 300
(config-if)#no nd-reachable-time
```

© 2026 IP Infusion Inc. Proprietary

no debug all

Use this command to disable all ARP and the neighbor debugging.

Command Syntax

```
no debug all
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#no debug all
```

© 2026 IP Infusion Inc. Proprietary

show arp

Use this command to display ARP entry information.

Command Syntax

```
show arp (| (A.B.C.D | detail | static) (| vrf (all | VRFNAME | default)))
show arp IFNAME (| vrf (all | VRFNAME | default))
show arp summary (|IFNAME) (| vrf (all | VRFNAME | default))
```

Parameters

detail
Display detailed information.

static
Display static ARP entries.

A.B.C.D
Specify the IP address of the ARP entry.

IFNAME
Specify the name of the interface.

all
All VRFs.

VRFNAME
VPN routing/forwarding instance name.

default
Default VRF.

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show arp

Flags: D - Static Adjacencies attached to down interface

IP ARP Table for context default
Total number of entries: 2
Address          Age      MAC Address     Interface  State
10.12.18.1       00:00:18 44e4.d982.274b  eth0       REACHABLE
10.10.10.20      00:02:33 a8b1.d433.4140  xe1        STALE

#show arp summary vrf default

IP ARP Table - Adjacency Summary

  Resolved   : 2
```

```
Incomplete : 0
Unknown    : 0
Total      : 2
```

Table 64 explains the show commmand output fields.

Table 64. show arp output details

Entry	Description
Address	ARP entry's IP address in the device.
Age	The number of minutes before the ARP entry was refreshed. If this value reaches the ARP aging period, the entry is removed from the table. Static entries do not age out.
MAC Address	Physical address of the host.
Interface	Logical address to connect the device over network.
State	INCOMPLETE – first ARP request sent, send ARP request. REACHABLE – normal expiration reset use counter. STALE – still usable; needs verification reset use counter; change state to delay. DELAY – schedule ARP request; needs verification reset use counter. PROBE – sending ARP request reset use counter. FAILED – no response received send ARP request. NOARP – normal expiration; never verified reset use counter. PERMANENT – never expires; never verified reset use counter. **Note:** INCOMPLETE/FAILED state entries are only cleared if the total number of ARP entries is more than 2048. It might take up to 60 seconds to clear the entries. This is an optimization to improve ND performance.

Network devices are considered adjacent if they can reach each other with a single hop. The summary command shows the count of the state of devices that are adjacencies.

Table 65 explains the show command output fields.

Table 65. show arp summary output details

Field	Description
Resolved	Count of working/known adjacencies.
Incomplete	Count of yet to be established adjacencies.
unknown	Count of adjacencies not currently in ARP table.
Total	Total count of all adjacencies.

© 2026 IP Infusion Inc. Proprietary

show debugging ip arp

Use this command to display debugging information for ARP.

Command Syntax

```
show debugging ip arp
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show debugging ip arp
ND IP Debugging status:
ND ip arp event debugging is off
ND ip arp packet debugging is off
```

© 2026 IP Infusion Inc. Proprietary

show debugging ipv6 nd

Use this command to display debugging information for the neighbor.

Command Syntax

```
show debugging ipv6 nd
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show debugging ipv6 nd
ND IPV6 Debugging status:
ND ipv6 event debugging is off
ND ipv6 packet debugging is off
```

© 2026 IP Infusion Inc. Proprietary

show ipv6 neighbors

Use this command to display the neighbor information.

Command Syntax

```
show ipv6 neighbors (| (X:X::X:X | detail | static) (| vrf (all | VRFNAME | default)))
show ipv6 neighbors IFNAME (| vrf (all | VRFNAME | default))
show ipv6 neighbors summary (| IFNAME) (| vrf (all | VRFNAME | default))
```

Parameters

detail
Show detail information of neighbor.

static
Static entry.

X:X::X:X
Specify the neighbor's IPv6 address.

IFNAME
Specify the name of the interface.

all
All VRFs.

VRFNAME
VPN routing/forwarding instance name.

default
Default VRF.

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ipv6 neighbors

R - Reachable, I - Incomplete, S - Stale, F - Failed, P - Probe,
D - Delay, PR - Permanent

Flags: (D) - Static neighbors attached to down interface

IPv6 Neighbor Table for context default
Total number of entries:2
Address        Age       MAC Address     Source  Interface      State
fe80::210:18ff:fe7f:f758
               00:43:04  0010.187f.f758  icmpv6  eth0             S
2000::5        00:55:25  0000.0001.0242  icmpv6  xe1              S

#show ipv6 neighbors summary
```

```
IPv6 Neighbors Table - Adjacency Summary

  Resolved   : 2
  Incomplete : 0
  Unknown    : 0
  Total      : 2
```

Table 66 shows the status codes displayed at the start of a route entry.

Table 66. status code output details

Status Code	Field	Description
R	Reachable	Normal expiration reset use counter.
I	Incomplete	First ARP request sent, send ARP request.
S	Stale	Still usable; needs verification reset use counter; change state to delay.
F	Failed	ARP requests response not received.
P	Probe	ARP request reset use counter.
D	Delay	Schedule ARP request; needs verification reset use counter.
PR	Permanent	Never expires; never verified reset use counter.

Table 67 explains the show command output fields.

Table 67. show ipv6 neighbors output details

Field	Description
Address	ARP entry's IP address in the device.
Age	The number of minutes before the ARP entry was refreshed. If this value reaches the ARP aging period, the entry is removed from the table. Static entries do not age out.
MAC Address	Physical address of the host.
Source	ARP request source in the interface.
Interface	Logical address to connect the device over network.
State	ARP request state that is being handled by the IPv6 neighbor session.

Table 68 explains the show command output fields.

Table 68. show arp summary output details

Field	Description
Resolved	Count of working/known adjacencies.

Table 68. show arp summary output details (continued)

Field	Description
Incomplete	Count of yet to be established adjacencies.
unknown	Count of adjacencies not currently in ARP table.
Total	Total count of all adjacencies.

© 2026 IP Infusion Inc. Proprietary

OPEN SHORTEST PATH FIRST CONFIGURATION

OSPFv2 2130
Enable OSPF on an Interface 2130
Set Priority 2134
Area Border Router 2140
Redistribute Routes into OSPF 2144
Cost 2146
Path MTU for IPv4 and IPv6 2152
Virtual Links 2155
OSPF Authentication 2160
Multiple OSPF Instances 2165
Multiple OSPF Instances on Same Subnet 2177
Multi-Area Adjacency Configuration 2180
LSA Throttling 2184
Loop-Free Alternate Fast Reroute 2196
Loop-Free Alternate (LFA) ECMP PATH 2206
SNMP Support for Multiple Instance of OSPF Configuration 2214
OSPF Stub Router Advertisement 2236
Multi-Area Redundant Adjacency Configuration 2243
Overview 2243
Topology 2244
Configuration 2245
Sample show running-config Output 2246
Validation 2248
Implementation Example 2254
OSPFv3 2256
Enable OSPFv3 on an Interface 2257
Set Priority 2260
Area Border Router 2264
Redistribute Routes into OSPFv3 2269
Cost 2274
Virtual Links 2283
Multiple Instances 2289
Not-So-Stubby Area 2303
NSSA with the Summary Address Option 2307
NSSA with the Translator Role Option 2311
Link LSA Suppression 2316
Originate Type-7 LSAs and Translate to Type-5 2318

© 2026 IP Infusion Inc. Proprietary

Summarize Inter-Area and External Routes 2322
Distribute List 2327
OSPFv3 Authentication 2331
OSPFv3 Graceful Restart Configuration 2336
Topology 2336
Configure R1 for Graceful Restart 2336
Configure R2 as Restart Helper 2337
Remove Capability Restart Configuration from R1 2337
Remove Helper Configuration from R2 2337
Validation 2338
OSPF Sham-link for VPN Sites Configuration 2342
Topology 2342
CE1 Configuration 2342
PE1 Configuration 2343
PE2 Configuration 2345
CE2 Configuration 2347
Validation 2348
OSPF TE-Metric Extension 2352
Configuration 2352
OSPFv2 LDP Remote Loop-Free Alternate (RLFA) 2358
RLFA Terminology 2358
Establishing RLFA Tunnel 2360
Assumptions and Limitations 2361
Topology 2361

© 2026 IP Infusion Inc. Proprietary

OSPFv2

This section 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 144. Basic OSPF Topology

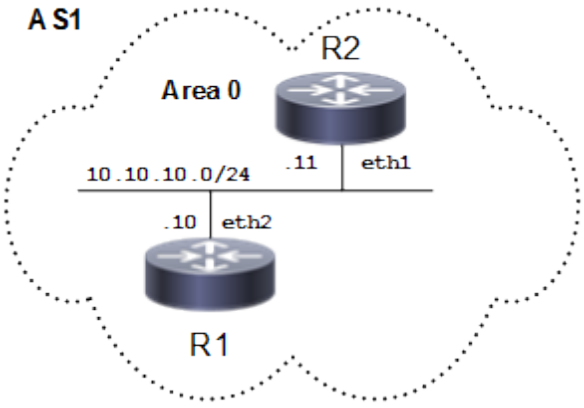

R1

`#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).
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

R1

The following provides the R1 validation:

```
#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


#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

The following provides the R2 validation:

```
#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 145. Set OSPF Priority

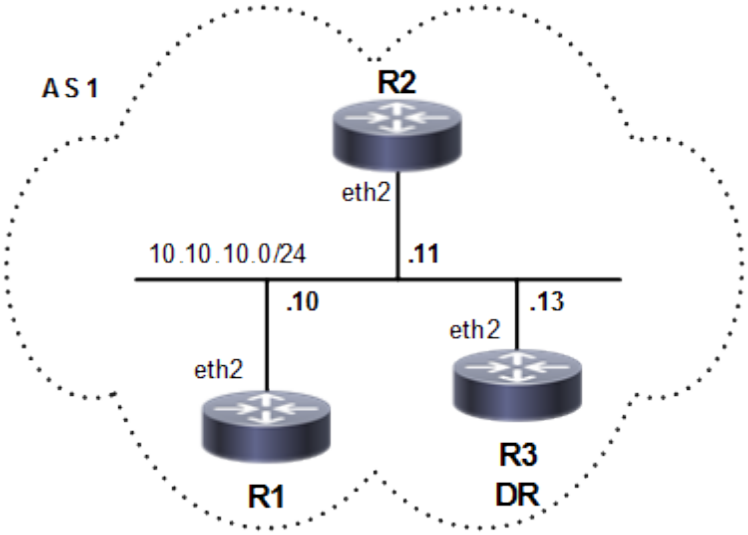

R3

`#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.
`(config-router)#commit`	Commit the candidate configuration to the running 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 10.10.10.0/24 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)#commit`	Commit the candidate configuration to the running configuration.

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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

R1

The following provides the R1 validation:

```
#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

!
line con 0
 login
line vty 0 39
 login
!
end
```

R2

The following provides the R2 validation:

```
#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
!
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

!
line con 0
 login
line vty 0 39
 login
!
end

#show 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#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 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

The following provides the R3 validation:

```
#show running-config
!
no service password-encryption
!
hostname R3
```

```
!
logging monitor 7
!
ip vrf management
!
ip domain-lookup
spanning-tree mode provider-rstp
data-center-bridging enable
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

!
line con 0
 login
line vty 0 39
 login
!
end

#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
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


#show 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

© 2026 IP Infusion Inc. Proprietary

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 146. OSPF ABR Topology

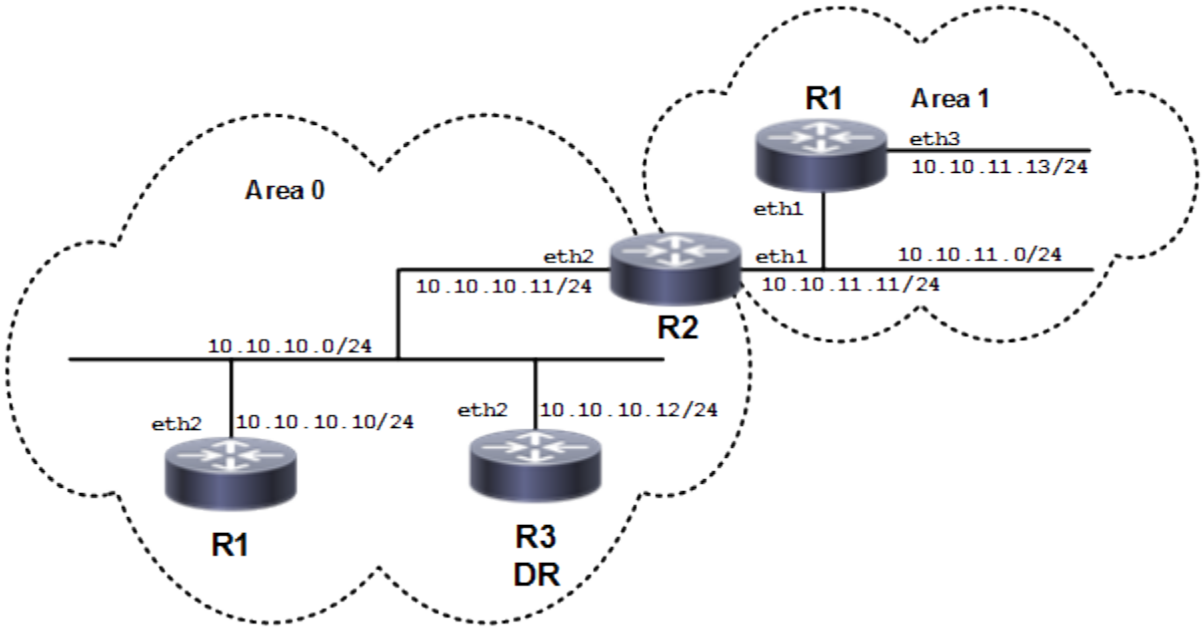

Configuration

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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

Validation

R2

The following provides the R2 validation:

```
#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

!
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
```

© 2026 IP Infusion Inc. Proprietary

```
 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 147. Redistribute Routes

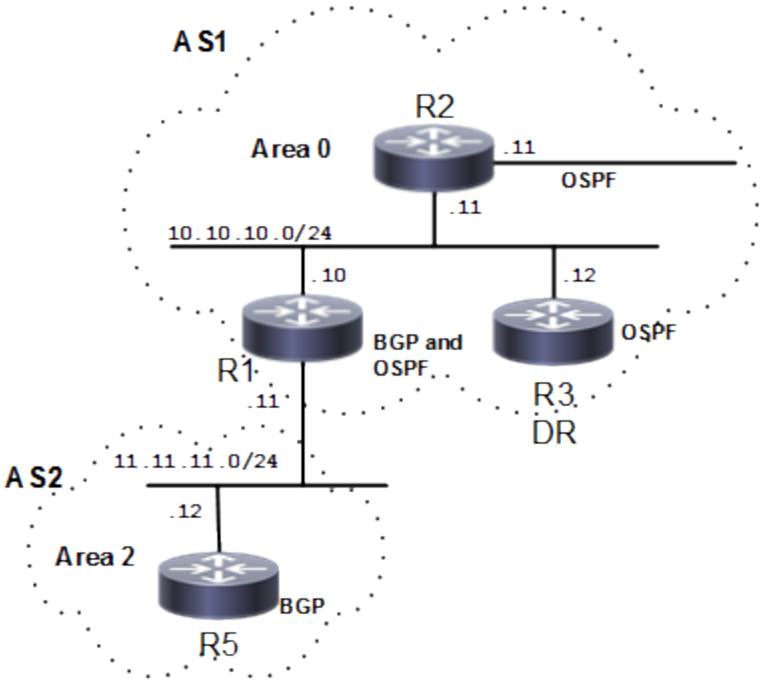

R1

`#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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

```
#show ip ospf route
```

© 2026 IP Infusion Inc. Proprietary

```
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
```

© 2026 IP Infusion Inc. Proprietary

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 148. Configure Cost Topology

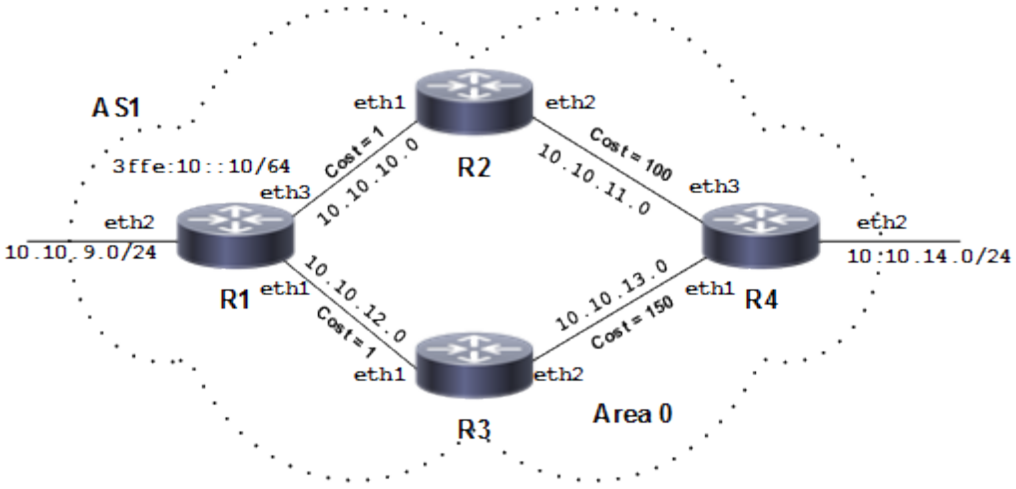

Configuration

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`	Define interfaces on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#network 10.10.10.0/24 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)#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).
(config-router)#commit	Commit the candidate configuration to the running configuration.

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	Define the interfaces on which OSPF runs, and associate the area ID (0) with the interface.
(config-router)#network 10.10.11.0/24 area 0	Define the interfaces on which OSPF runs, and associate the area ID (0) with the interface.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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	Define interfaces on which OSPF runs, and associate the area ID (0) with the interface.
(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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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	Define interfaces on which OSPF runs, and associate the area ID 0) with the interface.

© 2026 IP Infusion Inc. Proprietary

`(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.
`(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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

R1

The following provides the R1 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

#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

The following provides the R2 validation:

```
#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

The following provides the R3 validation:

```
#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
  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

The following provides the R4 validation:

```
#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 149. Path MTU Topology

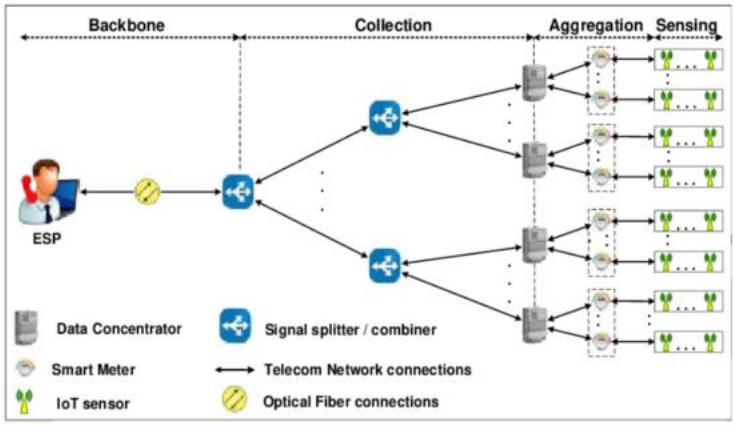

R1

`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

© 2026 IP Infusion Inc. Proprietary

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
`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 150. Virtual Links Topology

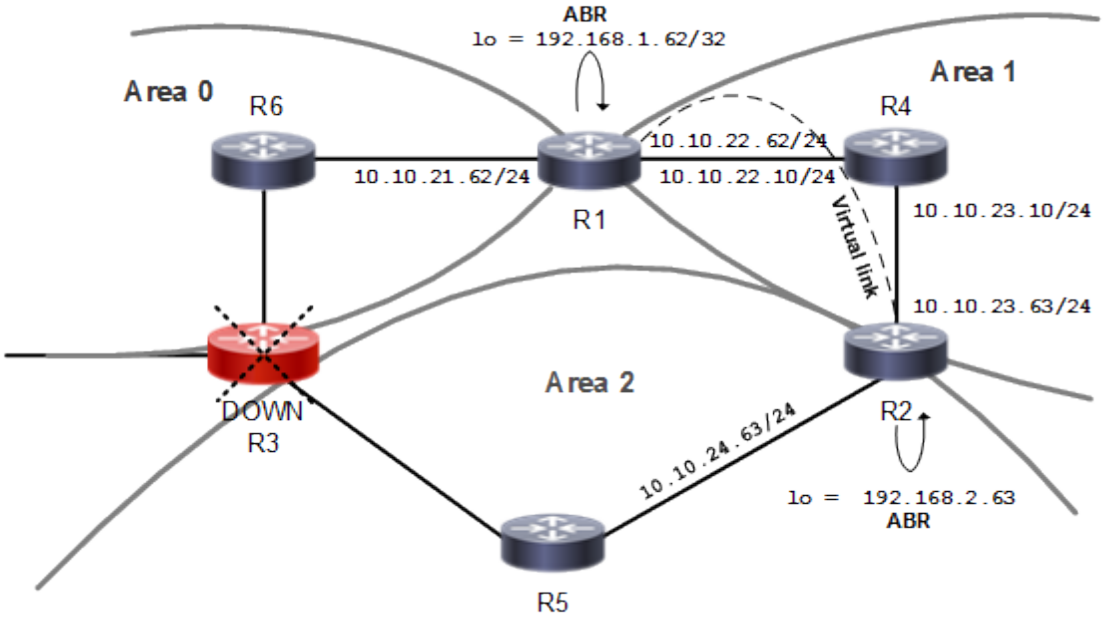

R1

`#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 secondary`	Configure the IP address of the interface loopback.
`(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`	Define interfaces on which OSPF runs, and associate the area IDs `(0)` with the interface.
`(config-router)#network 10.10.22.0/24 area 1`	Define interfaces on which OSPF runs, and

	associate the area IDs (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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

R2

`(config)#interface lo`	Specify loopback as the interface you want to configure.
`(config-if)#ip address 192.168.2.63/32 secondary`	Configure the IP address of the interface loopback.
`(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`	Define interfaces on which OSPF runs, and associate the area IDs `(1)` with the interface.
`(config-router)#network 10.10.24.0/24 area 2`	Define interfaces on which OSPF runs, and associate the area IDs `(2)` with the interface.
`(config-router)#network 192.168.2.63/32 area 2`	Define interfaces on which OSPF runs, and associate the area IDs `(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`.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

The following provides the validation for virtual links configuration:

```
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
 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 eth2 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 151. OSPF Authentication Topology

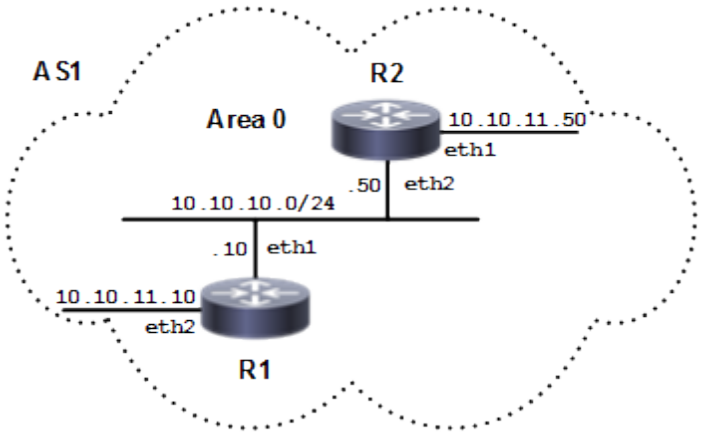

R1

`#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`	Define interfaces on which OSPF runs, and associate the area ID `(0)` with the interface (area ID 0 specifies the backbone area).
`(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 0x293da85becc67703`	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 0x3a24102d157b1618`	Specify an OSPF authentication password (`test`) for the neighboring routers.
`(config-if)#commit`	Commit the candidate configuration to the running 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 identifying the routing process.
`(config-router)#network 10.10.10.0/24 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)#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 eth2`	Enter interface mode.
`(config-if)# ip ospf authentication-key 0x293da85becc67703`	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 message-digest-key 1 md5 0x293da85becc67703`	Specify an OSPF authentication password `test` for the neighboring routers.
`(config-if)#commit`	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

Validation

R1

The following provides the R1 validation:

```
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

!
line con 0
 login
line vty 0 39
 login
!
end


R1#sh ip ospf neighbor
```

© 2026 IP Infusion Inc. Proprietary

```
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

The following provides the R1 validation:

```
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

!
line con 0
 login
line vty 0 39
 login
!
end

R2#sh ip ospf neighbor
```

© 2026 IP Infusion Inc. Proprietary

```
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 152. Multiple OSPF Instances

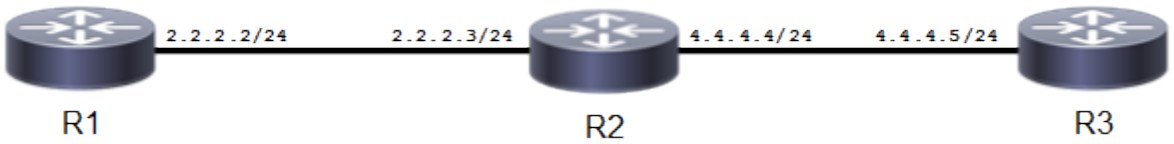

.

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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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.

© 2026 IP Infusion Inc. Proprietary

(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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

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.
(config-router)#commit	Commit the candidate configuration to the running configuration.

Validation

R1

The following provides the R1 validation:

```
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

!
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

The following provides the R2 validation:

```
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

!
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

The following provides the R3 validation:

```
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

!
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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

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
```

© 2026 IP Infusion Inc. Proprietary

```
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
```

© 2026 IP Infusion Inc. Proprietary

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 153. Multiple Instances on the Same Subnet

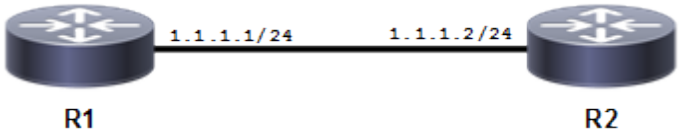

Configuration

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)#commit`	Commit the candidate configuration to the running configuration.
`(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.

© 2026 IP Infusion Inc. Proprietary

`(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)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit Router mode, and return to Configure mode.

Validation

R1

The following provides the R1 validation:

```
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

The following provides the R2 validation:

```
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 154. One Subnet with Multiple OSPF Areas

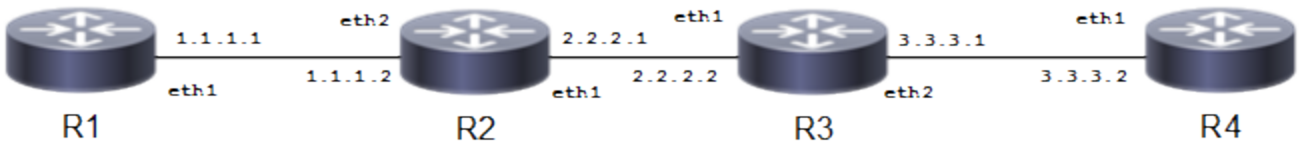

Configuration

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)#commit`	Commit the candidate configuration to the running configuration.
`(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)#commit`	Commit the candidate configuration to the running configuration.
`(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.

© 2026 IP Infusion Inc. Proprietary

`(config-if)#ip ospf 1 multi-area 0.0.0.1 neighbor 2.2.2.2`	Configure multi area adjacency.
`(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 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)#commit`	Commit the candidate configuration to the running configuration.
`(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)#commit`	Commit the candidate configuration to the running configuration.
`(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)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	Exit Router mode, and return to Configure mode.

Validation

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 155. Basic OSPF Topology

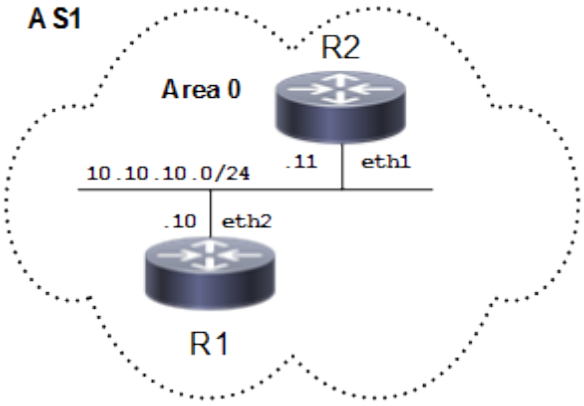

R1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Specify the interface loopback to configure.
`(config-if)#ip address 1.1.1.1/32 secondary`	Configure the ip address (1.1.1.1) of the 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)#logging level ospf 7`	Enable logging level ospf globally.
`(config)#commit`	Commit the candidate configuration to the running configuration.
`(config)#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)#commit`	Commit the candidate configuration to the running configuration.

Validation

R1

The following provides the R1 validation:

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

The following provides the R2 validation:

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 156. Basic OSPF Topology

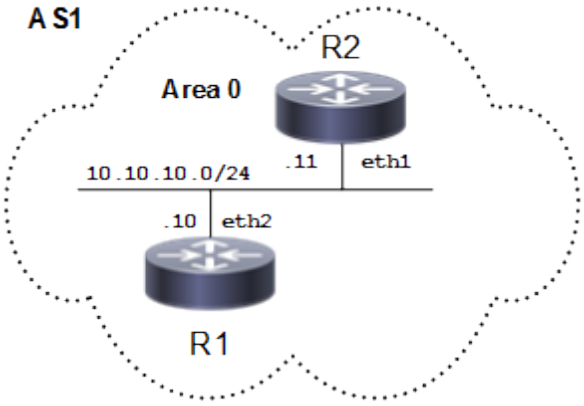

Configuration

R1

`#configure terminal`	Enter configure mode.
`(config)#interface lo`	Specify the interface loopback to configure.
`(config-if)#ip address 1.1.1.1/32 secondary`	Configure the ip address (1.1.1.1) of the 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)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#end`	Exit router mode

© 2026 IP Infusion Inc. Proprietary

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.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

R1

The following provides the R1 validation:

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

The following provides the R2 validation:

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.

Topology

The diagram shows the configuration required to enable the OSPF LFA feature.

Figure 157. Basic OSPF-LFA Topology

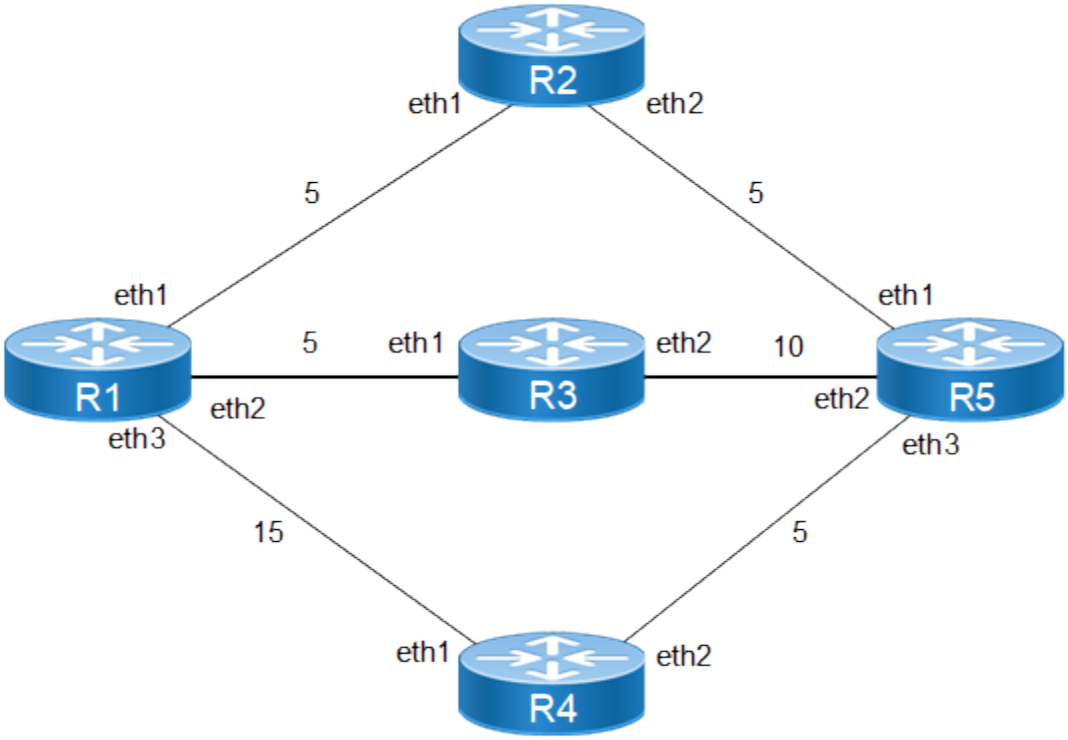

R1

`#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).
`(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)#commit`	Commit the candidate configuration to the running configuration.
`(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

© 2026 IP Infusion Inc. Proprietary

	(area ID 0 specifies the backbone area).
(config-router)#commit	Commit the candidate configuration to the running configuration.
(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)#commit	Commit the candidate configuration to the running configuration.
(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).

© 2026 IP Infusion Inc. Proprietary

`(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)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#end`	Exit router mode.

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)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#end`	Exit router mode.

Validation

R1

The following provides the R1 validation:

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
(config-router)#commit
```

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
```

© 2026 IP Infusion Inc. Proprietary

```
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
```

© 2026 IP Infusion Inc. Proprietary

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)#commit`	Commit the candidate configuration to the running configuration.
`(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 [110/2 0] 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:

© 2026 IP Infusion Inc. Proprietary

`(config)#interface eth2`	Enter interface mode.
`(config-if)#shutdown`	Shutdown the interface
`(config-if)#commit`	Commit the candidate configuration to the running configuration.
`(config-if)#exit`	Exit interface mode
`(config-if)#exit`	Exit interface mode.

```
#show ip route fast-rerouteCodes: 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 [110/2 0] via 10.1.1.2, eth1, 00:02:27
                    [FRR-NH] via 30.1.1.2, eth3
```

© 2026 IP Infusion Inc. Proprietary

Loop-Free Alternate (LFA) ECMP PATH

This section contains configurations for OSPF LFA ECMP which provides LFA/alternate path from primary ECMP path set or non-primary/non-ECMP path set which improve convergence after a primary path failure occur in network.

Overview

With ECMP, a prefix has multiple primary paths to forward traffic. When a particular primary path fails, the other primary paths are not guaranteed to provide protection against the failure scenario. As part of LFA ECMP, alternate paths are determined for each primary path separately. The selected alternate path can be either one of the primary path from the set of ECMP or a loop-free non-ECMP if available.

In OSPF, by default the LFA algorithm tries to find loop free node protecting alternate from the set of existing primary next-hops. If no loop free node-protecting alternate is available, the LFA algorithm tries to find link-protecting alternate from the set of existing primary next-hops. If no loop-free node-protecting and link-protecting alternate is available, then the LFA algorithm should select a loop-free link-protecting from the non-ECMP next-hops.

Note: RIB errors occur when the hardware profile cannot handle LFA-ECMP. The error prompts users to alter or reverse the configurations.

Topology

Figure 158. OSPF-LFA_ECMP

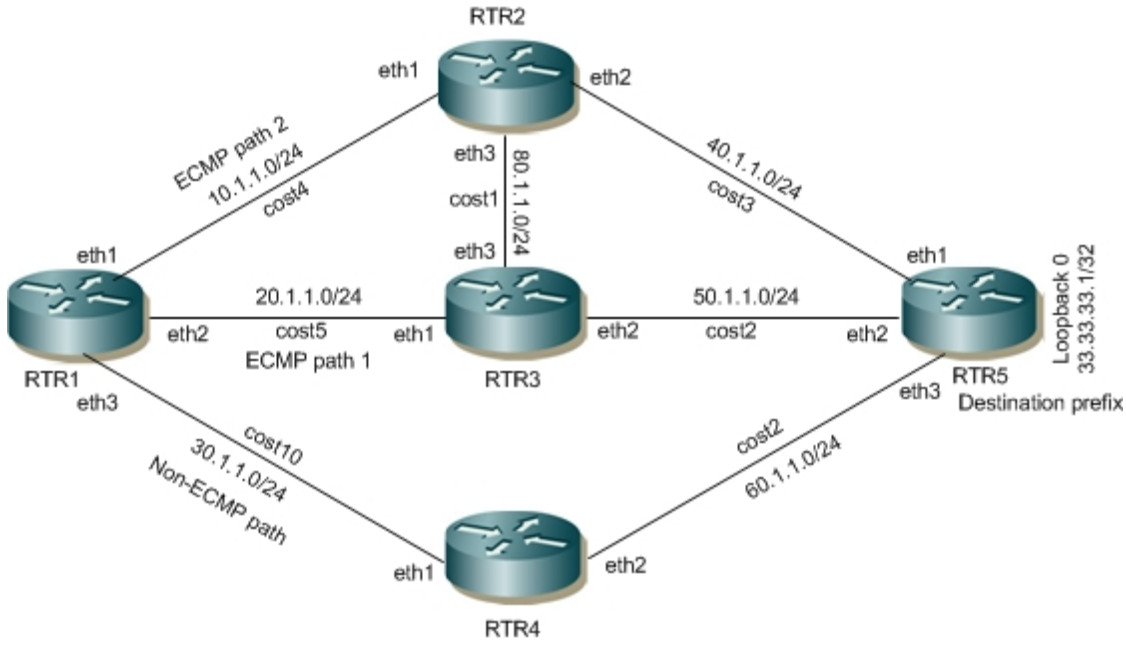

Configuring OSPF LFA ECMP

Configuration Part 1: with default LFA configuration where primary path priority higher than Secondary-path (non-ECMP) and LFA selection happen within primary ECMP path

RTR1

#configure terminal	Enter configure mode.
(config)#interface eth1	Enter interface mode.
(config-if)#ip address 10.1.1.1/24	Configure the IP address of the interface.
(config-if)#ip ospf cost 4	Assign cost to interface
(config-if)#exit	Exit interface mode.
(config)#interface eth2	Enter interface mode.
(config-if)# ip address 20.1.1.1/24	A s si gn IP ad dr e s s .
(config-if)#ip ospf cost 5	Assign cost to interface
(config-if)#exit	Exit interface mode.
(config)#interface eth3	Enter interface mode.
(config-if)# ip address 30.1.1.1/24	A s si gn IP ad dr e s s .
(config-if)#ip ospf cost 10	Assign cost to interface
(config-if)#exit	Exit interface mode.
(config)#router ospf 1	Configure the routing process and specify the Process ID (1).
(config-router)#ospf router-id 1.1.1.1	Configure router-id as 1.1.1.1
(config-router)# bfd all-interfaces	Enable BFD over ospf for all ospf enabled interfaces
(config-router)#network 10.1.1.0/24 area 0	C o nf i g ure OSPF ne tw ork f or ar ea 0 .
(config-router)#network 20.1.1.0/24 area 0	C o nf i g ure OSPF ne tw ork f or ar ea 0 .
(config-router)#network 30.1.1.0/24 area 0	C o nf i g ure OSPF ne tw ork f or ar ea 0 .
(config-router)#fast-reroute keep-all-paths	Configure LFA-FRR to calculate the available backup path.
(config-router)#commit	Commit the candidate configuration to the running configuration.
(config-router)#end	Exit router mode.

RTR2

#configure terminal	Enter configure 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)#ip ospf cost 4	Assign cost to interface
(config-if)#exit	Exit interface mode.

© 2026 IP Infusion Inc. Proprietary

(config)#interface eth2	Enter interface mode.
(config-if)# ip address 40.1.1.1/24	A s si gn IP ad dr e s s .
(config-if)#ip ospf cost 3	Assign cost to interface
(config-if)#exit	Exit interface mode.
(config)#interface eth3	Enter interface mode.
(config-if)# ip address 80.1.1.1/24	A s si gn IP ad dr e s s .
(config-if)#ip ospf cost 1	Assign cost to interface
(config-if)#exit	Exit interface mode.
(config)#router ospf 1	Configure the routing process and specify the Process ID (1).
(config-router)#ospf router-id 2.2.2.2	Configure router-id.
(config-router)# bfd all-interfaces	Enable BFD over ospf for all ospf enabled interfaces
(config-router)#network 10.1.1.0/24 area 0	C o nf i g ure OSPF ne tw ork f or ar ea 0 .
(config-router)#network 40.1.1.0/24 area 0	C o nf i g ure OSPF ne tw ork f or ar ea 0 .
(config-router)#network 80.1.1.0/24 area 0	C o nf i g ure OSPF ne tw ork f or ar ea 0 .
(config-router)#commit	Commit the candidate configuration to the running configuration.
(config-router)#end	Exit router mode.

RTR3

#configure terminal	Enter configure mode.
(config)#interface eth1	Enter interface mode.
(config-if)#ip address 20.1.1.2/24	Configure the IP address of the interface.
(config-if)#ip ospf cost 5	Assign cost to interface
(config-if)#exit	Exit interface mode.
(config)#interface eth2	Enter interface mode.
(config-if)# ip address 50.1.1.1/24	A s si gn IP ad dr e s s .
(config-if)#ip ospf cost 2	Assign cost to interface
(config-if)#exit	Exit interface mode.
(config)#interface eth3	Enter interface mode.
(config-if)# ip address 80.1.1.2/24	A s si gn IP ad dr e s s .
(config-if)#ip ospf cost 1	Assign cost to interface
(config-if)#exit	Exit interface mode.
(config)#router ospf 1	Configure the routing process and specify the Process ID (1).
(config-router)#ospf router-id 3.3.3.3	Configure router-id.
(config-router)# bfd all-interfaces	Enable BFD over ospf for all ospf enabled interfaces

© 2026 IP Infusion Inc. Proprietary

(config-router)#network 20.1.1.0/24 area 0	C o nf i g ure OSPF ne tw ork f or ar ea 0 .
(config-router)#network 50.1.1.0/24 area 0	C o nf i g ure OSPF ne tw ork f or ar ea 0 .
(config-router)#network 80.1.1.0/24 area 0	C o nf i g ure OSPF ne tw ork f or ar ea 0 .
(config-router)#commit	Commit the candidate configuration to the running configuration.
(config-router)#end	Exit router mode.

RTR4

#configure terminal	Enter configure mode.
(config)#interface eth1	Enter interface mode.
(config-if)#ip address 30.1.1.2/24	Configure the IP address of the interface.
(config-if)#ip ospf cost 10	Assign cost to interface
(config-if)#exit	Exit interface mode.
(config)#interface eth2	Enter interface mode.
(config-if)# ip address 60.1.1.1/24	A s si gn IP ad dr e s s .
(config-if)#ip ospf cost 2	Assign cost to interface
(config-if)#exit	Exit interface mode.
(config)#router ospf 1	Configure the routing process and specify the Process ID (1).
(config-router)#ospf router-id 4.4.4.4	Configure router-id.
(config-router)# bfd all-interfaces	Enable BFD over ospf for all ospf enabled interfaces
(config-router)#network 30.1.1.0/24 area 0	C o nf i g ure OSPF ne tw ork f or ar ea 0 .
(config-router)#network 60.1.1.0/24 area 0	C o nf i g ure OSPF ne tw ork f or ar ea 0 .
(config-router)#commit	Commit the candidate configuration to the running configuration.
(config-router)#end	Exit router mode.

RTR5

#configure terminal	Enter configure mode.
(config)#interface lo	Enter interface mode.
(config-if)# ip address 33.33.33.1/32 secondary	Configure the IP address of the interface loopback.
(config-if)#exit	Exit interface mode.
(config)#interface eth1	Enter interface mode.
(config-if)#ip address 40.1.1.2/24	Configure the IP address of the interface.
(config-if)#ip ospf cost 3	Assign cost to interface
(config-if)#exit	Exit interface mode.

© 2026 IP Infusion Inc. Proprietary

(config)#interface eth2	Enter interface mode.
(config-if)# ip address 50.1.1.2/24	A s si gn IP ad dr e s s .
(config-if)#ip ospf cost 2	Assign cost to interface
(config-if)#exit	Exit interface mode.
(config)#interface eth3	Enter interface mode.
(config-if)# ip address 60.1.1.2/24	A s si gn IP ad dr e s s .
(config-if)#ip ospf cost 2	Assign cost to interface
(config-if)#exit	Exit interface mode.
(config)#router ospf 1	Configure the routing process and specify the Process ID (1).
(config-router)#ospf router-id 5.5.5.5	Configure router-id.
(config-router)# bfd all-interfaces	Enable BFD over ospf for all ospf enabled interfaces
(config-router)#network 40.1.1.0/24 area 0	C o nf i g ure OSPF ne tw ork f or ar ea 0 .
(config-router)#network 50.1.1.0/24 area 0	C o nf i g ure OSPF ne tw ork f or ar ea 0 .
(config-router)#network 60.1.1.0/24 area 0	C o nf i g ure OSPF ne tw ork f or ar ea 0 .
(config-router)#network 33.33.33.1/32 area 0	C o nf i g ure OSPF ne tw ork f or ar ea 0 .
(config-router)#commit	Commit the candidate configuration to the running configuration.
(config-router)#end	Exit router mode.

Validation

Validation Part 1:LFA selected from primary ECMP path set

RTR1

The following provides the RTR1 validation:

```
#show ip ospf route fast-reroute

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
       OSPF LFA attributes:
       P - Primary, SP - Secondary-Path, ID - Interface Disjoint,
       NP - Node Protecting, BID - Broadcast Interface Disjoint

O  33.33.33.1/32 [17] 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: [17] ,P ,NP
                      via 20.1.1.2, eth2, Area 0.0.0.0
                        Backup path:
                            via 10.1.1.2, eth1, Area 0.0.0.0
                           Attributes: Metric: [17] ,P ,ID
O  40.1.1.0/24 [7] via 10.1.1.2, eth1, Area 0.0.0.0
                        Backup path:
                          via 30.1.1.2, eth3, Area 0.0.0.0
                         Attributes: Metric: [15] ,SP ,NP
O  50.1.1.0/24 [7] via 20.1.1.2, eth2, Area 0.0.0.0
```

```
                        Backup path:
                         via 10.1.1.2, eth1, Area 0.0.0.0
                        Attributes: Metric: [7] ,P ,ID
                    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: [7] ,P ,NP
O  60.1.1.0/24 [9] 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: [9] ,P ,NP
                    via 20.1.1.2, eth2, Area 0.0.0.0
                        Backup path:
                         via 10.1.1.2, eth1, Area 0.0.0.0
                        Attributes: Metric: [9] ,P ,ID
O  80.1.1.0/24 [5] 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: [6] ,SP ,NP

#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, T - FRR nhp,p - stale info
       * - candidate default

IP Route Table for VRF "default"
O        33.33.33.1/32 [110/17] via 20.1.1.2, eth2, 00:00:44
                       [FRR-NH] via 10.1.1.2, eth1

                       [110/17] via 10.1.1.2, eth1
                       [FRR-NH] via 20.1.1.2, eth2

O        40.1.1.0/24 [110/7] via 10.1.1.2, eth1, 00:01:46
                     [FRR-NH] via 30.1.1.2, eth3

O        50.1.1.0/24 [110/7] via 10.1.1.2, eth1, 00:01:34
                     [FRR-NH] via 20.1.1.2, eth2

                     [110/7] via 20.1.1.2, eth2
                     [FRR-NH] via 10.1.1.2, eth1

O        60.1.1.0/24 [110/9] via 20.1.1.2, eth2, 00:01:34
                     [FRR-NH] via 10.1.1.2, eth1

                     [110/9] via 10.1.1.2, eth1
                     [FRR-NH] via 20.1.1.2, eth2

O        80.1.1.0/24 [110/5] via 10.1.1.2, eth1, 00:01:46
                     [FRR-NH] via 20.1.1.2, eth2


#show ip ospf route 33.33.33.1

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
       OSPF LFA attributes:
       P - Primary, SP - Secondary-Path, ID - Interface Disjoint,
       NP - Node Protecting, BID - Broadcast Interface Disjoint

O  33.33.33.1/32 [17] via 10.1.1.2, eth1, Area 0.0.0.0
                      via 20.1.1.2, eth2, Area 0.0.0.0
```

```
#show ip route 33.33.33.1
Routing entry for 33.33.33.1/32
  Known via "ospf", distance 110, metric 17,  External Route Tag: 0, best
  Last update 00:00:40 ago
  * 20.1.1.2, via eth2
  * 10.1.1.2, via eth1
```

Configuration Part 2: with non-ECMP tiebreaker configured where secondary-path priority higher than primary (ECMP) path

Configure below configuration with config's shown in Part1:

RTR1

`#configure terminal`	Enter configure mode.
`(config)# router ospf 1`	E n t er R o ut er OSPF m o de .
`(config-router)#fast-reroute tie-break secondary-path index 5`	Configure LFA tiebreaker for LFA to be calculate from non-ecmp path set if available (this is user defined to decide the priority to select between ecmp/non-ecmp set)
`(config-router)#commit`	Commit the candidate configuration to the running configuration.
`(config-router)#exit`	E x it R o ut er OSPF m ode a nd r e t urn to C o nf i g ure m o de .

Validation Part 2: LFA selected from non-ecmp path for each primary ecmp path

RTR1

The following provides the RTR1 validation

```
#show ip ospf route fast-reroute

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
       OSPF LFA attributes:
       P - Primary, SP - Secondary-Path, ID - Interface Disjoint,
       NP - Node Protecting, BID - Broadcast Interface Disjoint

O  33.33.33.1/32 [17] via 10.1.1.2, eth1, Area 0.0.0.0
                        Backup path:
                            via 30.1.1.2, eth3, Area 0.0.0.0
                           Attributes: Metric: [22] ,SP ,NP
                      via 20.1.1.2, eth2, Area 0.0.0.0
                        Backup path:
                            via 30.1.1.2, eth3, Area 0.0.0.0
                           Attributes: Metric: [22] ,SP ,NP
O  40.1.1.0/24 [7] via 10.1.1.2, eth1, Area 0.0.0.0
                        Backup path:
                          via 30.1.1.2, eth3, Area 0.0.0.0
                         Attributes: Metric: [15] ,SP ,NP
O  50.1.1.0/24 [7] via 20.1.1.2, eth2, Area 0.0.0.0
                        Backup path:
                          via 30.1.1.2, eth3, Area 0.0.0.0
                         Attributes: Metric: [14] ,SP ,NP
                    via 10.1.1.2, eth1, Area 0.0.0.0
```

```
                        Backup path:
                          via 30.1.1.2, eth3, Area 0.0.0.0
                         Attributes: Metric: [14] ,SP ,NP
O  60.1.1.0/24 [9] via 10.1.1.2, eth1, Area 0.0.0.0
                        Backup path:
                          via 30.1.1.2, eth3, Area 0.0.0.0
                         Attributes: Metric: [12] ,SP ,NP
                   via 20.1.1.2, eth2, Area 0.0.0.0
                        Backup path:
                          via 30.1.1.2, eth3, Area 0.0.0.0
                         Attributes: Metric: [12] ,SP ,NP
O  80.1.1.0/24 [5] 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: [6] ,SP ,NP

#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 , T - FRR nhp,p - stale info
       * - candidate default

IP Route Table for VRF "default"
O        33.33.33.1/32 [110/17] via 20.1.1.2, eth2, 00:00:36
                       [FRR-NH] via 30.1.1.2, eth3

                       [110/17] via 10.1.1.2, eth1
                       [FRR-NH] via 30.1.1.2, eth3

O        40.1.1.0/24 [110/7] via 10.1.1.2, eth1, 00:01:38
                     [FRR-NH] via 30.1.1.2, eth3

O        50.1.1.0/24 [110/7] via 10.1.1.2, eth1, 00:01:26
                     [FRR-NH] via 30.1.1.2, eth3

                     [110/7] via 20.1.1.2, eth2
                     [FRR-NH] via 30.1.1.2, eth3

O        60.1.1.0/24 [110/9] via 20.1.1.2, eth2, 00:01:26
                     [FRR-NH] via 30.1.1.2, eth3

                     [110/9] via 10.1.1.2, eth1
                     [FRR-NH] via 30.1.1.2, eth3

O        80.1.1.0/24 [110/5] via 10.1.1.2, eth1, 00:01:38
```

[FRR-NH] via 20.1.1.2, eth2

Configuring OSPF LFA ECMP with Load Balancing

For OSPF LFA ECMP with load balancing, there are two options:

- Option 1: Have more than two links for ECMP, where one path is used for FRR and the other paths are used for load balancing.
- Option 2: Configure OSPF LFA tie-breaker with secondary path, where the secondary path is used for FRR and the other path becomes the primary path for load balancing.

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 159. SNMP OSPF instance

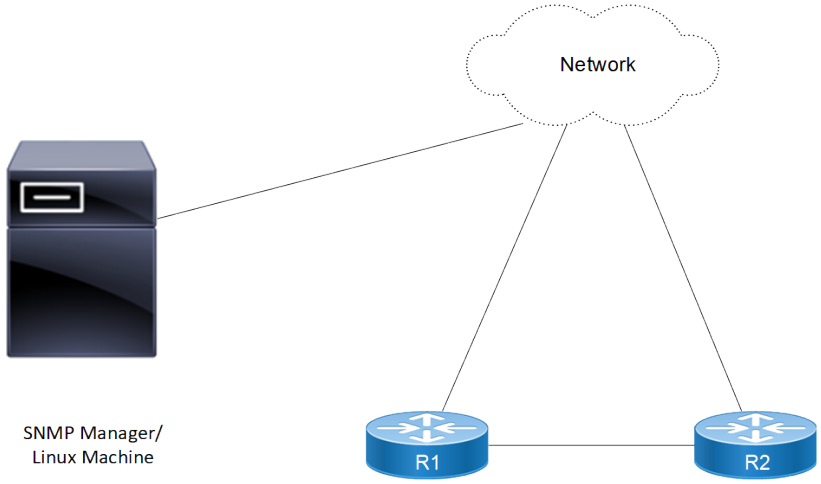

Enable SNMP and create SNMP Context & Group for OSPF Instances on a Router with SNMPv2

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

The following provides the R1 validation:

```
R1#sh running-config
!
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

The following provides the R2 validation:

```
R2#sh running-config
!
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
```

© 2026 IP Infusion Inc. Proprietary

```
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 and 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

The following provides the R1 validation:

```
R1#sh running-config
!
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
```

© 2026 IP Infusion Inc. Proprietary

```
R1#
```

R2

The following provides the R2 validation

```
R2#  sh run
!
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
```

© 2026 IP Infusion Inc. Proprietary

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 160. SNMP OSPF multiple instance

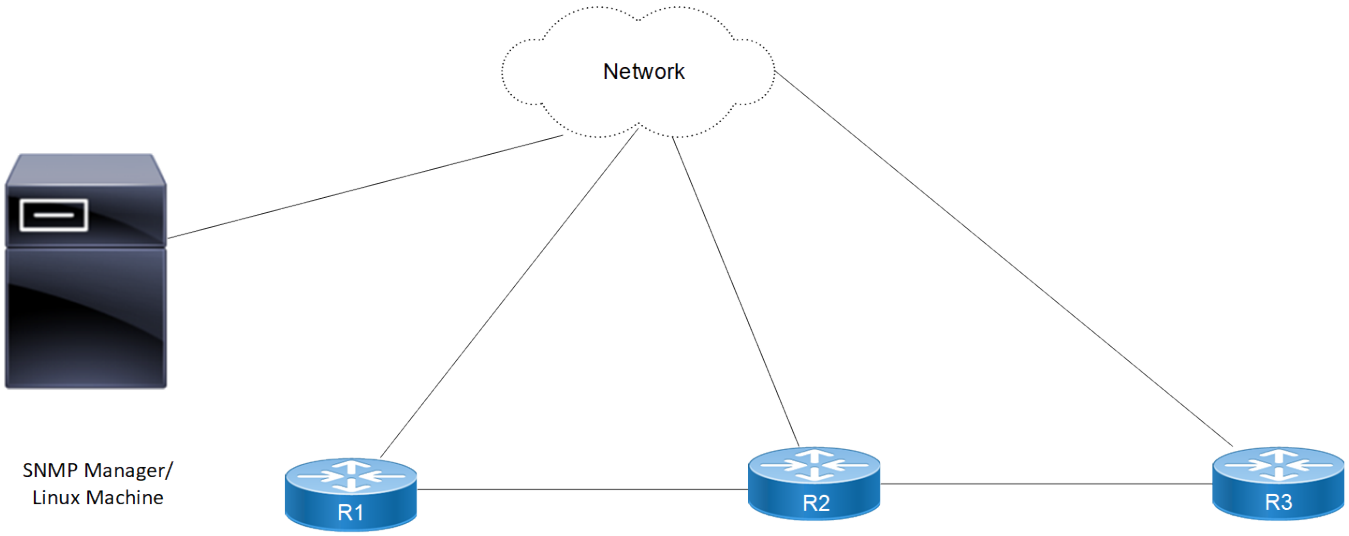

Configurations

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

© 2026 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.

© 2026 IP Infusion Inc. Proprietary

(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

© 2026 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 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

The following provides the R1 validation:

```
R1#show run
!
! Software version: EC_AS7326-56X-OcNOS-5.1.194-DC-MPLS-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

The following provides the R2 validation:

```
R2(config)#show run
!
! Software version: EC_AS5912-54X-OcNOS-5.1.194-SP-MPLS-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

The following provides the R3 validation:

```
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
```

© 2026 IP Infusion Inc. Proprietary

```
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
```

© 2026 IP Infusion Inc. Proprietary

```
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
```

© 2026 IP Infusion Inc. Proprietary

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 161. OSPF Stub Router Advertisement Topology

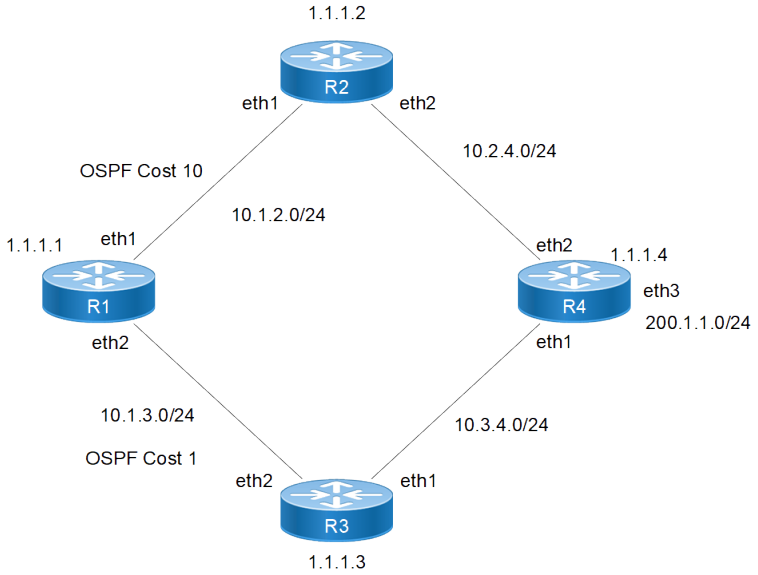

Configuration

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

© 2026 IP Infusion Inc. Proprietary

(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).

© 2026 IP Infusion Inc. Proprietary

`(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

The following provides the R3 validation:

```
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#
```

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.

© 2026 IP Infusion Inc. Proprietary

Prerequisites

The router must be configured with OSPFv2 with Multi-Area support. For more information, refer to .

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.

The following topology illustrates the Multi-Area Adjacency routing information.

Figure 162. OSPFv2 Multi-Area Adjacency with Multiple Interfaces

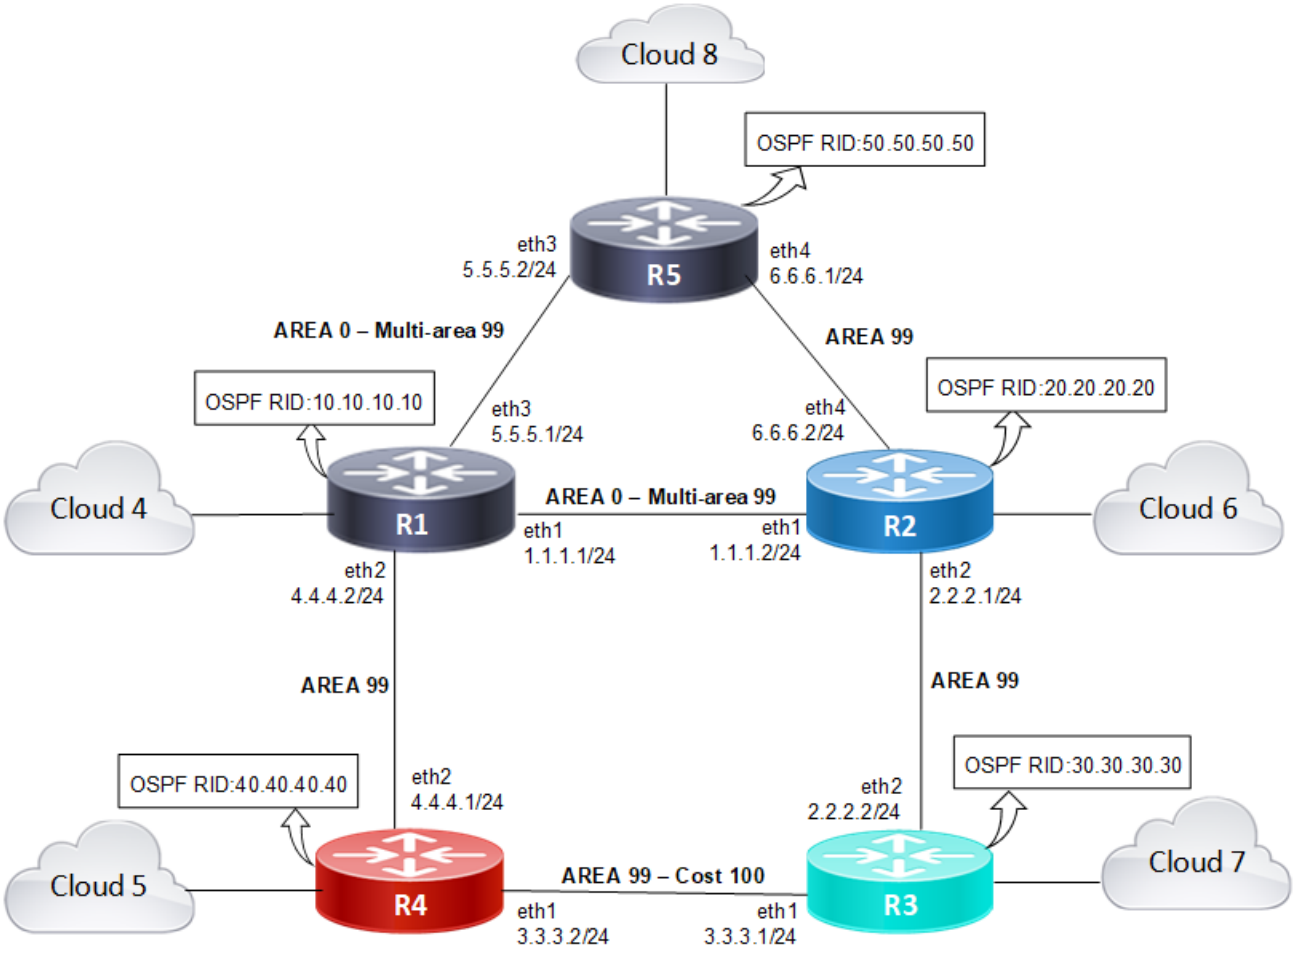

The Multi-Area support changes this default behavior and permits that once Multi-Area is enabled on `eth1` Area 0 for Multi-Area 99 in R1, the routing communication between R1 and R3 is via R1-R2-R3 as `eth1` Area 99 link cost is high.

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.

© 2026 IP Infusion Inc. Proprietary

Configuration

The following configuration enables the Multi-Area feature on multiple interfaces in the OcNOS devices.

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
   ```

© 2026 IP Infusion Inc. Proprietary

2. Configure an OSPF instance with an instance ID on R1, R2, R3, R4, and R5. Refer to theTopology (page 2244) 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
```

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
```

© 2026 IP Infusion Inc. Proprietary

```
 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
```

© 2026 IP Infusion Inc. Proprietary

```
  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
```

© 2026 IP Infusion Inc. Proprietary

```
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
```

© 2026 IP Infusion Inc. Proprietary

```
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,
```

© 2026 IP Infusion Inc. Proprietary

```
       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 (page 2471)

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.

© 2026 IP Infusion Inc. Proprietary

OSPFv3

This section contains basic OSPFv3 configuration examples.

© 2026 IP Infusion Inc. Proprietary

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 163. Basic OSPFv3 Topology

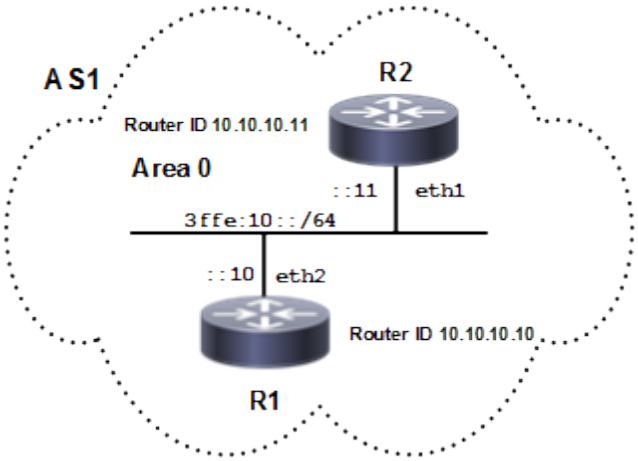

Configuration

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.

© 2026 IP Infusion Inc. Proprietary

`(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 164. OSPFv3 Set Priority

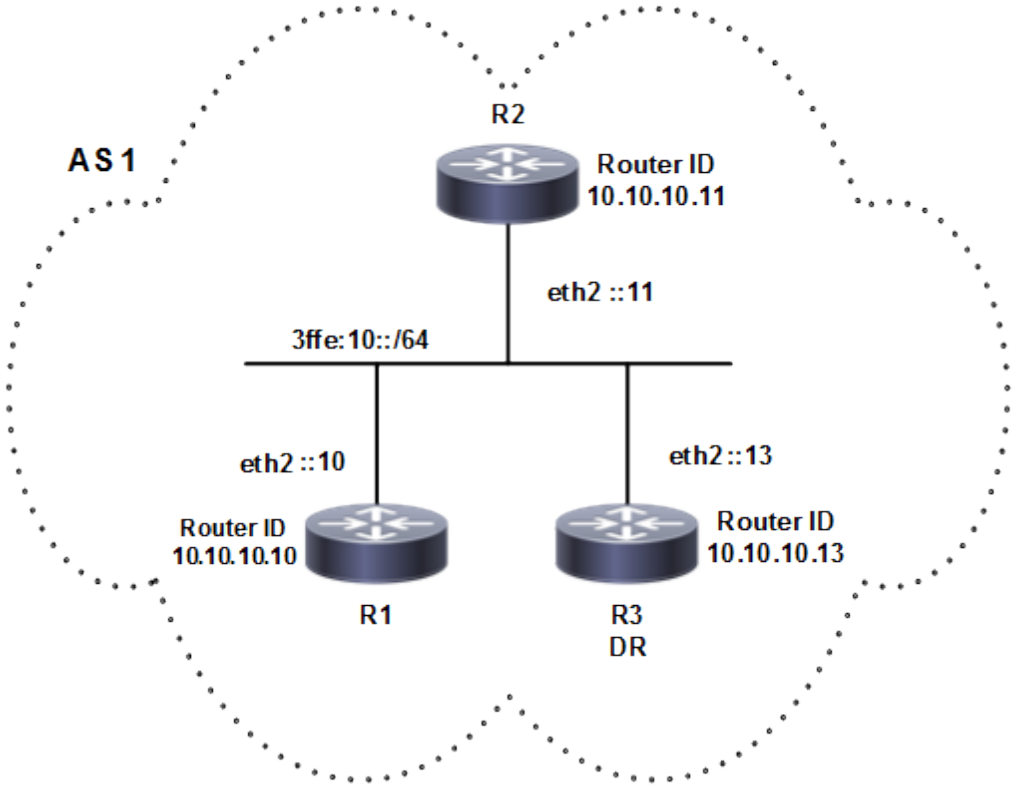

Configuration

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

© 2026 IP Infusion Inc. Proprietary

	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 165. OSPFv3 Area Border Router

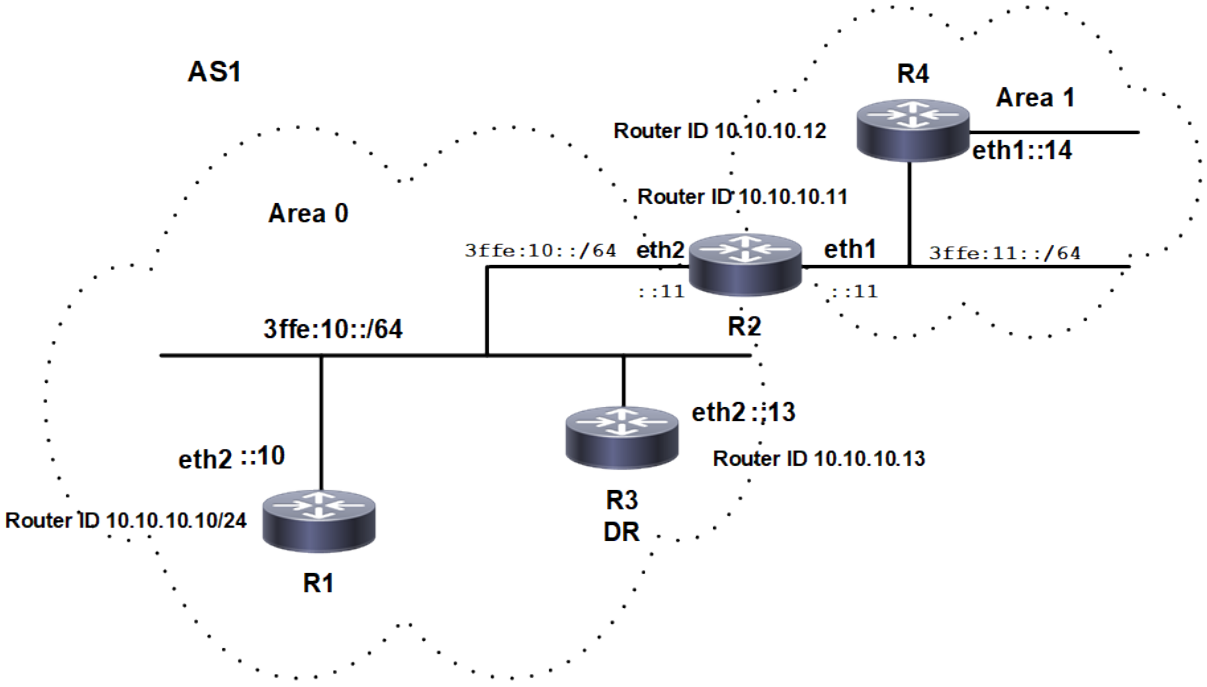

Configuration

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)`.

© 2026 IP Infusion Inc. Proprietary

`(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#
```

© 2026 IP Infusion Inc. Proprietary

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#
```

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 166. OSPFv3 Redistribute Routes

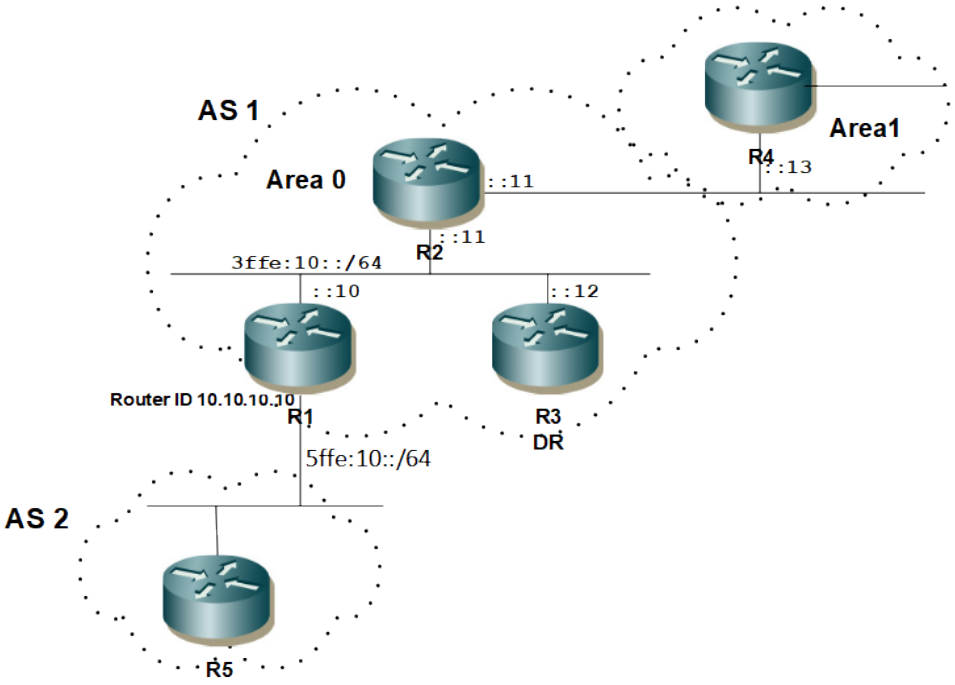

Configuration

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.

© 2026 IP Infusion Inc. Proprietary

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#
```

© 2026 IP Infusion Inc. Proprietary

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 167. Configure Cost OSPFv3

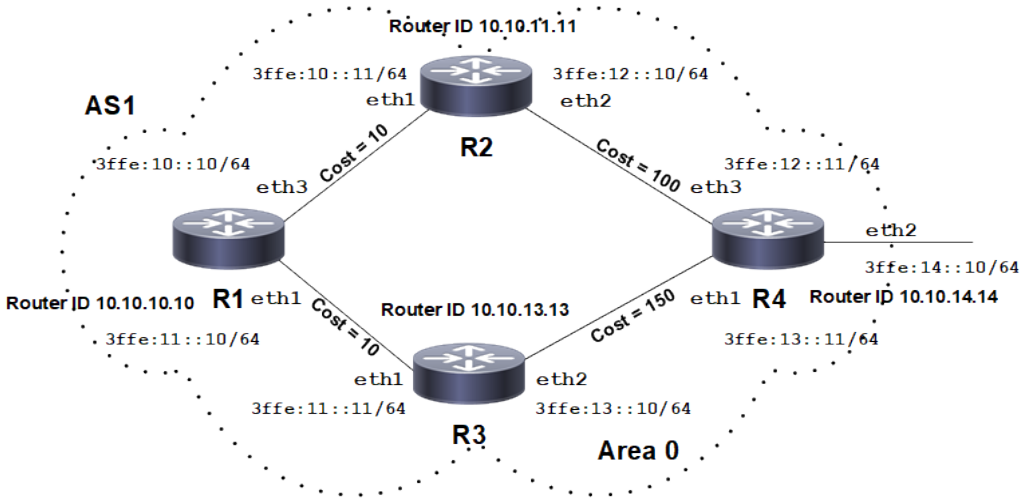

Configuration

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)`.

© 2026 IP Infusion Inc. Proprietary

`(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.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

```
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 168. OSPFv3 Virtual Links

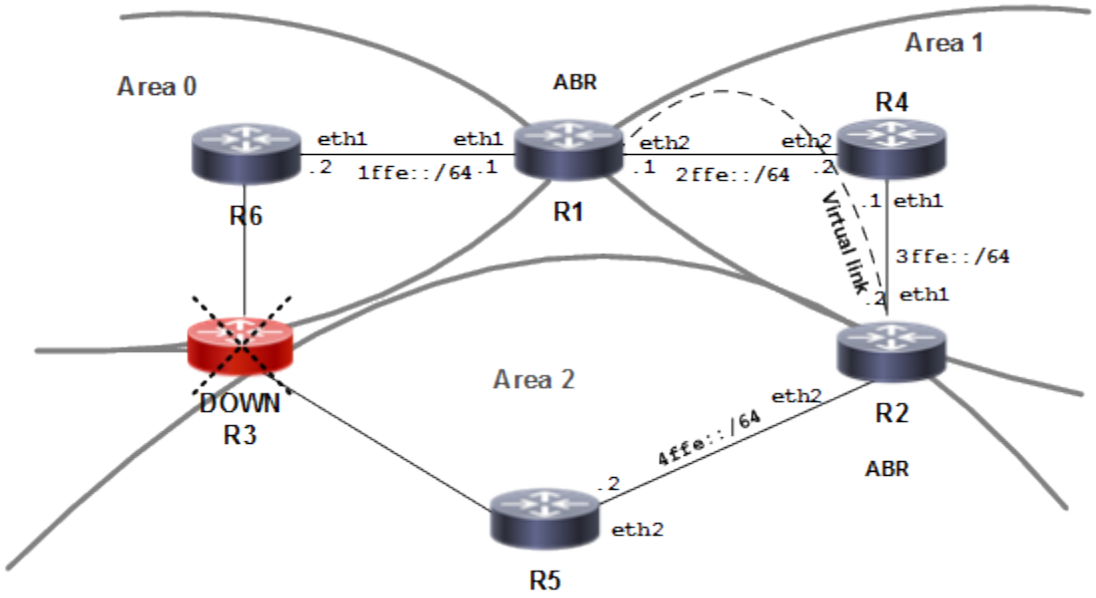

Configuration

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.

© 2026 IP Infusion Inc. Proprietary

`(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

© 2026 IP Infusion Inc. Proprietary

	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

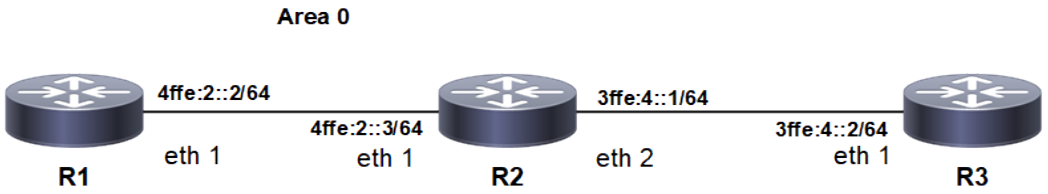

Enable Multiple OSPFv3 Instances on a Router Based on Tags

In this example, routers R1, R2, and R3 are in Area 0, and all run OSPFv3.

Configuration

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.

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

`(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.

Configuration

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)
```

© 2026 IP Infusion Inc. Proprietary

```
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.

Configuration

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.

Configuration

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
```

© 2026 IP Infusion Inc. Proprietary

```
       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

Configuration

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.

© 2026 IP Infusion Inc. Proprietary

`(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 169. Translating Type 7 LSAs into Type 5 LSAs

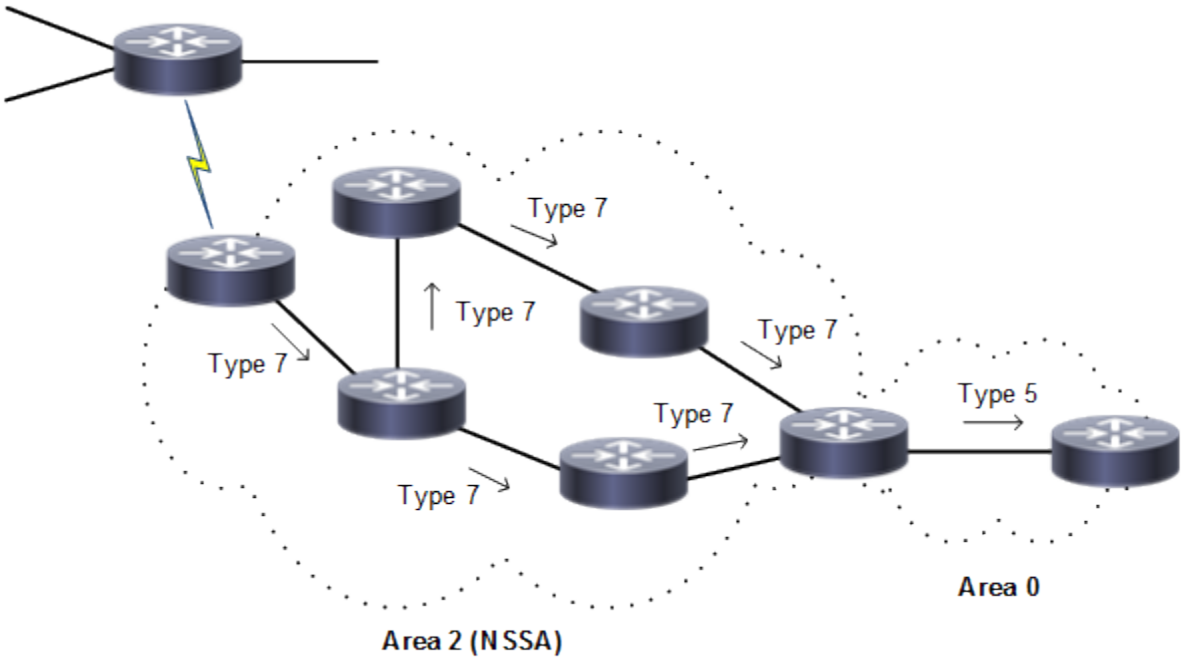

NSSA with Route Option

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 170, 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 170. NSSA with Route Options

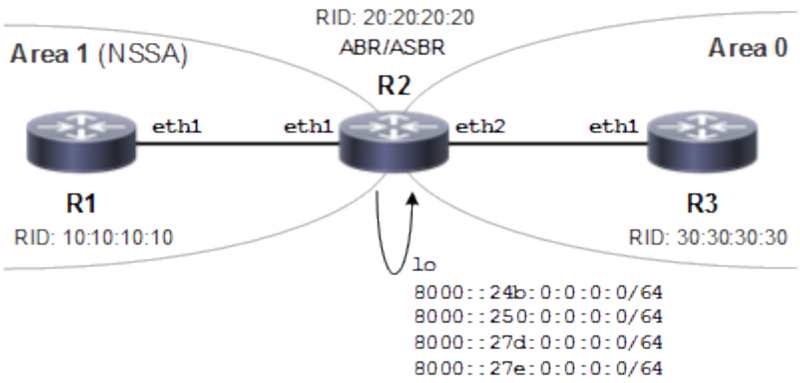

Configuration

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

© 2026 IP Infusion Inc. Proprietary

(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.
(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 171 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 171. Using the summary-address Option

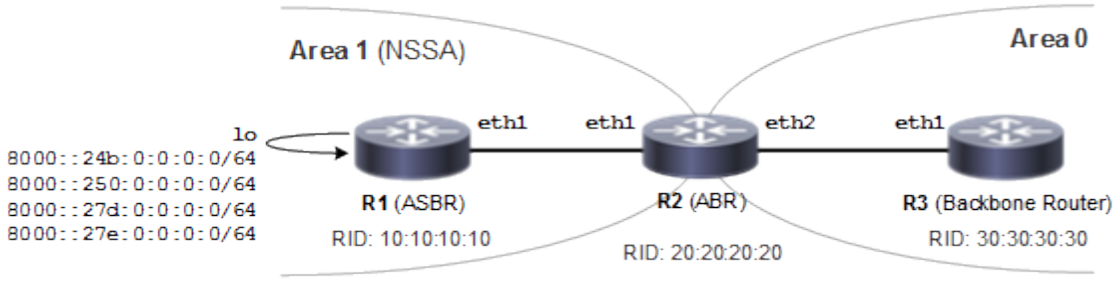

Configuration

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)

© 2026 IP Infusion Inc. Proprietary

`(config-router)#area 1 nssa`	Configure the area as NSSA.
`(config-router)#redistribute connected`	Redistribute the configured loopback network into OSPFv3 NSSA. 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

© 2026 IP Infusion Inc. Proprietary

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 below:

- 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 172. Using the translator-role Option

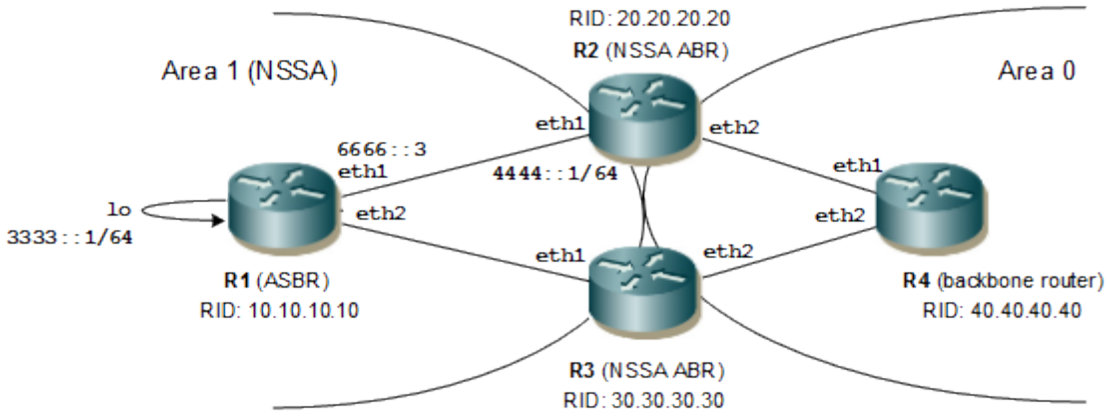

Configuration

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).

© 2026 IP Infusion Inc. Proprietary

`(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 173. LSA Suppression

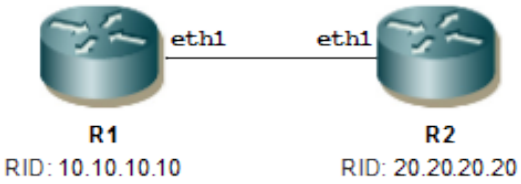

Configuration

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.

© 2026 IP Infusion Inc. Proprietary

`(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 174 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 174. Originate Type-7 LSAs and Translate to Type-5 under Address Family IPv4

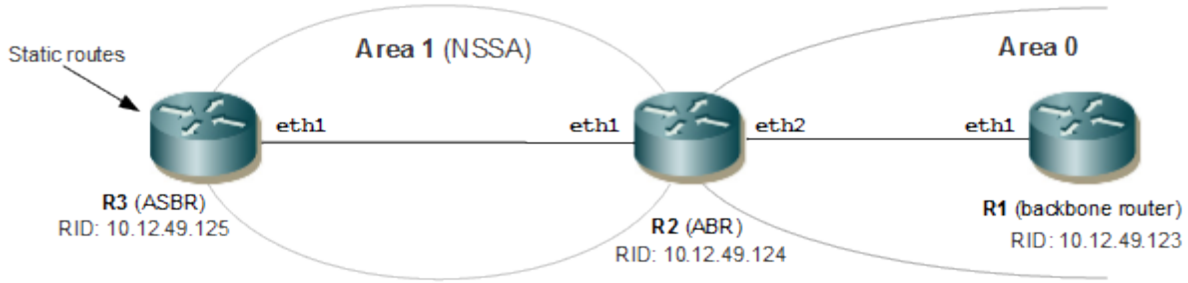

Configuration

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

© 2026 IP Infusion Inc. Proprietary

	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
`(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
```

© 2026 IP Infusion Inc. Proprietary

```
  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 4

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
```

© 2026 IP Infusion Inc. Proprietary

Summarize Inter-Area and External Routes

Figure 175 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 175. Enabling Intra-Area and External Route Summarization

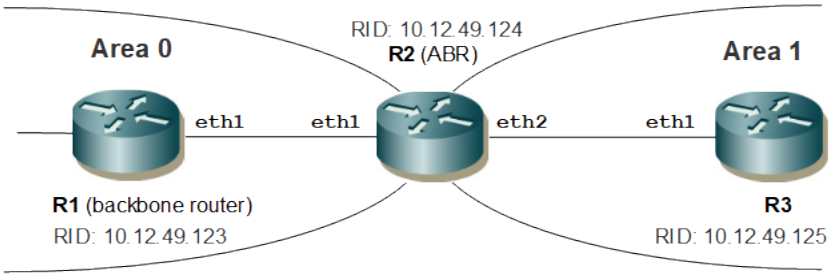

Configuration

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

© 2026 IP Infusion Inc. Proprietary

`(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
```

© 2026 IP Infusion Inc. Proprietary

```
  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#
```

© 2026 IP Infusion Inc. Proprietary

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 (-|-|-|-)
```

© 2026 IP Infusion Inc. Proprietary

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

Below figure shows the configuration to illustrate the distribute-list support for OSPFv3

Figure 176. Basic Topology for Distribute-list

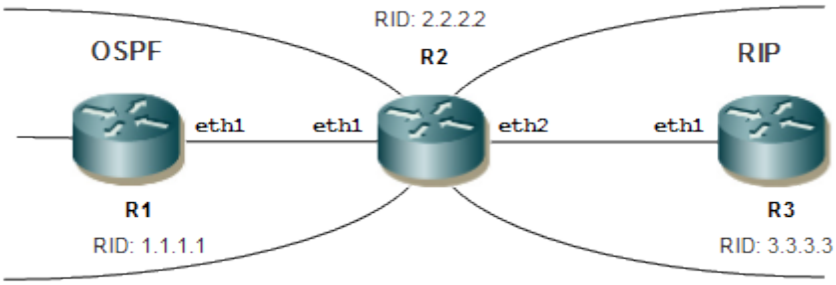

Configuration

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.

© 2026 IP Infusion Inc. Proprietary

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
`(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

© 2026 IP Infusion Inc. Proprietary

(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
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
```

© 2026 IP Infusion Inc. Proprietary

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.

Notes:

- You must explicitly specify a Router ID for the OSPFv3 process to be activated.
- Supported authentication algorithms are: MD5, SHA1, SHA256.

Topology

Figure 177. OSPFv3 Authentication

Configuration

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

© 2026 IP Infusion Inc. Proprietary

	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 3777217A25432A4676397924422645294840 4D6351655468	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
```

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 178. OSPFv3 Graceful Restart

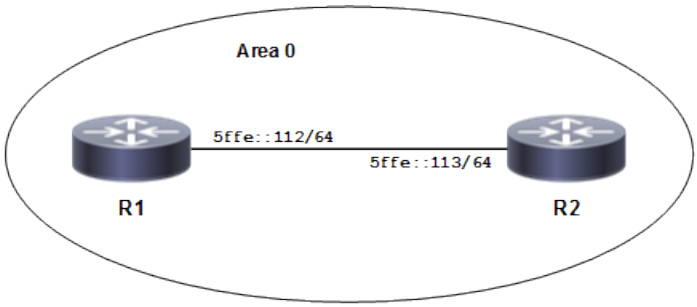

Configure R1 for Graceful Restart

`#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

© 2026 IP Infusion Inc. Proprietary

	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.

© 2026 IP Infusion Inc. Proprietary

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)
```

© 2026 IP Infusion Inc. Proprietary

OSPF Sham-link for VPN Sites Configuration

This feature is to ensure that the OSPF Client sites that share a backdoor link can communicate over the MPLS VPN backbone and participate in VPN services.

Suppose that there are two sites in same OSPF area and each of them is attached to a different PE router, and there is also an intra-area ospf backdoor link connecting the two sites. There will be routes between sites that go through the PE routers, but these routes will appear to be inter area routes, and OSPF will consider them less preferable than the intra-area routes through the backdoor link. To make a route through the backbone appear to be an intra-area route, it is necessary to make it appear as if there is an intra-area link connecting the two PE routers. Sham-links are those links routes the intra-area routes through the backbone.

Topology

Figure 179. OSPF SHAM-LINK FOR VPN SITES

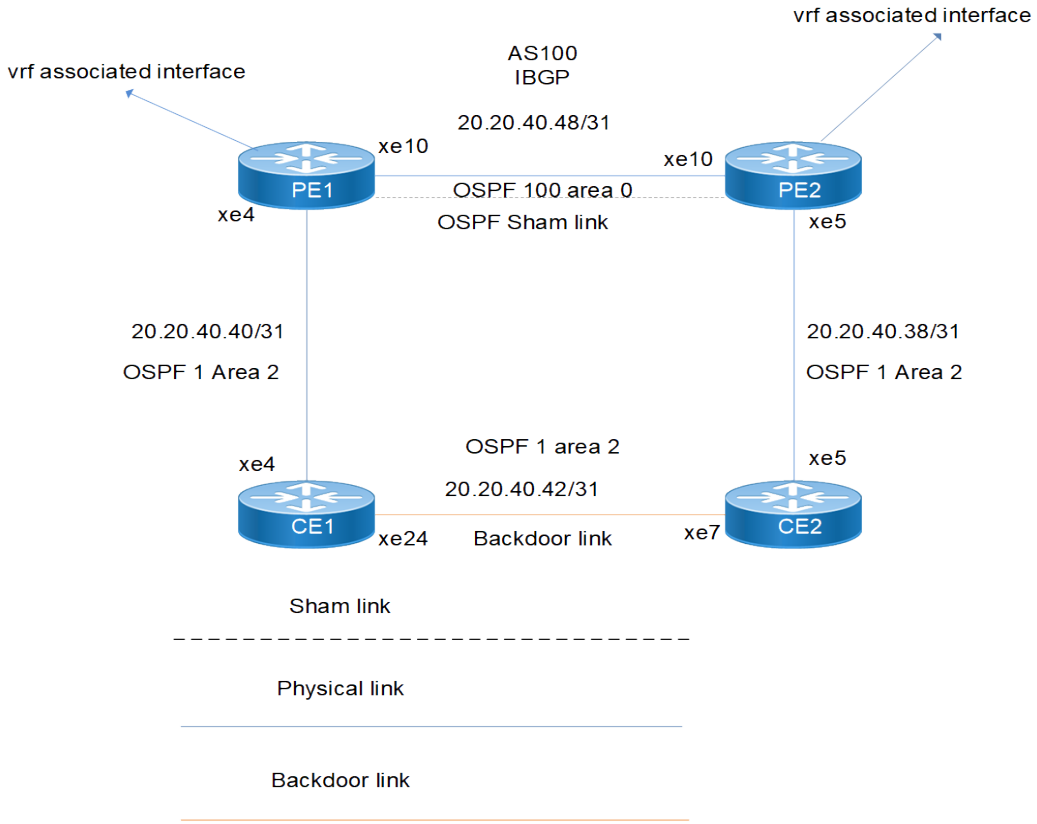

CE1 Configuration

`#configure terminal`	Enter the Configure mode.
`(config)#interface lo`	Enter interface mode
`(config-if)#ip address 1.1.1.1/32 secondary`	Configure the IP address of the interface loopback

(config-if)#exit	Exit interface mode
(config)# commit	Committing the configuration to apply in running configuration
(config)#interface xe4	Enter interface mode
(config-if)#ip address 20.20.40.40/31	Configure the IP address on the interface
(config-if)# description to_pe1	Adding Description to interface
(config-if)#exit	Exit interface mode.
(config)# commit	Committing the configuration to apply in running configuration
(config)#interface xe24	Enter interface mode
(config-if)#ip address 20.20.40.42/31	Configure the IP address on the interface
(config-if)# description ospf backdoor	Adding Description to interface
(config-if)# ip ospf cost 10	Set the OSPF cost of this link
(config-if)#exit	Exit interface mode.
(config)# commit	Committing the configuration to apply in running configuration
(config)#router ospf 1	Configure the routing process OSPF with process id
(config-router)#network 1.1.1.1/32 area 2	Define the interface on which OSPF runs, and associate the area ID with the interface.
(config-router)#network 20.20.40.40/31 area 2	Define the interface on which OSPF runs, and associate the area ID with the interface.
(config-router)#network 20.20.40.42/31 area 2	Define the interface on which OSPF runs, and associate the area ID with the interface.
(config-router)#exit	Exit from router ospf mode
(config)# commit	Committing the configuration to apply in running configuration

PE1 Configuration

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 2.2.2.2/32	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)# commit	Committing the configuration to apply in running configuration
(config)#ip vrf vrf1	Create vrf1
(config-vrf)# description vrf1	Adding description to vrf
(config-vrf)#rd 100:1	Specify the route distinguisher in the VRF
(config-vrf)#route-target both 100:1	Specify the import & export route target

© 2026 IP Infusion Inc. Proprietary

`(config)# commit`	Committing the configuration to apply in running configuration
`(config)# router ldp`	Configure Router LDP instance
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)#interface xe4`	Enter interface mode
`(config-if)# description to_ce1`	Adding description to interface
`(config-if)# ip vrf forwarding vrf1`	Associate the interface to vrf1
`(config-if)#ip address 20.20.40.41/31`	Configure the IP address on the interface
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe10`	Enter interface mode
`(config-if)# description to_pe2`	Adding description to interface
`(config-if)#ip address 20.20.40.48/31`	Configure the IP address on the interface
`(config-if)#label-switching`	Enable label-switching on interface
`(config-if)# enable-ldp ipv4`	Enable LDP process on interface
`(config-if)#exit`	Exit interface mode.
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)# interface lo.vrf1`	Enter interface mode
`(config-if)# ip vrf forwarding vrf1`	Associate the interface to vrf1
`ip address 11.11.11.11/32 secondary`	Configure the IP address on the interface
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)#router ospf 1 vrf1`	Associate the ospf process with vrf1.
`(config-router)# redistribute bgp`	Redistribute BGP into OSPF.
`(config-router)#network 20.20.40.40/31 area 2`	Define the interface on which OSPF runs, and associate the area ID with the interface.
`(config-router)# area 0.0.0.2 sham-link 11.11.11.11 22.22.22.22 cost 5`	Configuring Sham-link between PE routers with cost 5.
`(config-router)#exit`	Exit from router ospf mode
`(config)#router ospf 100`	Configure the routing process OSPF with process id
`(config-router)#network 2.2.2.2/32 area 2`	Define the interface on which OSPF runs, and associate the area ID with the interface.
`(config-router)#network 20.20.40.48/31 area 0`	Define the interface on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#exit`	Exit from router ospf mode
`(config)# commit`	Committing the configuration to apply in running configuration

© 2026 IP Infusion Inc. Proprietary

(config)#router bgp 100	Enter Router BGP mode
(config-router)#neighbor 3.3.3.3 remote-	
as 100	Add loopback ip of ABR2 as neighbor with neighbor AS
(config-router)#neighbor 3.3.3.3 update- source 2.2.2.2	Update the source for that particular neighbor as loopback interface
(config-router)#address-family vpnv4 unicast	Enter into vpnv4 unicast address family
(config-router-af)#neighbor 3.3.3.3 activate	Activate the neighbor inside vpnv4 address family
(config-router-af)#exit-address-family	Exit from address family vpnv4
(config-router)# address-family ipv4 vrf vrf1	Enter into vrf address family
(config-router-af)# network 11.11.11.11/32	Advertise a route via iBGP connection.
(config-router-af)#redistribute ospf 1	Redistribute ospf routes into bgp
(config-router-af)#exit-address-family	Exit from address family vpnv4
(config-router)#exit	Exit from router BGP mode
(config)# commit	Committing the configuration to apply in running configuration
(config)#end	Exit from config mode into privilege mode

PE2 Configuration

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 3.3.3.3/32	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)# commit	Committing the configuration to apply in running configuration
(config)#ip vrf vrf1	Create vrf1
(config-vrf)# description vrf1	Adding description to vrf
(config-vrf)#rd 100:1	Specify the route distinguisher in the VRF
(config-vrf)#route-target both 100:1	Specify the import & export route target
(config)# commit	Committing the configuration to apply in running configuration
(config)# router ldp	Configure Router LDP instance
(config)# commit	Committing the configuration to apply in running configuration
(config)#interface xe5	Enter interface mode
(config-if)# description to_ce2	Adding description to interface
(config-if)# ip vrf forwarding vrf1	Associate the interface to vrf1

© 2026 IP Infusion Inc. Proprietary

`(config-if)# ip address 20.20.40.38/31`	Configure the IP address on the interface
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe10`	Enter interface mode
`(config-if)# description to_pe1`	Adding description to interface
`(config-if)#ip address 20.20.40.49/31`	Configure the IP address on the interface
`(config-if)#label-switching`	Enable label-switching on interface
`(config-if)# enable-ldp ipv4`	Enable LDP process on interface
`(config-if)#exit`	Exit interface mode.
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)# interface lo.vrf1`	Enter interface mode
`(config-if)# ip vrf forwarding vrf1`	Associate the interface to vrf1
`ip address 22.22.22.22/32 secondary`	Configure the IP address on the interface
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)#router ospf 1 vrf1`	Associate the ospf process with vrf1.
`(config-router)# redistribute bgp`	Redistribute BGP into OSPF.
`(config-router)#network 20.20.40.38/31 area 2`	Define the interface on which OSPF runs, and associate the area ID with the interface.
`(config-router)# area 0.0.0.2 sham-link 22.22.22.22 11.11.11.11 cost 5`	Configuring Sham-link between PE routers with cost 5.
`(config-router)#exit`	Exit from router ospf mode
`(config)#router ospf 100`	Configure the routing process OSPF with process id
`(config-router)#network 3.3.3.3/32 area 2`	Define the interface on which OSPF runs, and associate the area ID with the interface.
`(config-router)#network 20.20.40.48/31 area 0`	Define the interface on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#exit`	Exit from router ospf mode
`(config)# commit`	Committing the configuration to apply in running configuration
`(config)#router bgp 100`	Enter Router BGP mode
`(config-router)#neighbor 2.2.2.2 remote-as 100`	Add loopback ip of ABR2 as neighbor with neighbor AS
`(config-router)#neighbor 2.2.2.2 update- source 3.3.3.3`	Update the source for that particular neighbor as loopback interface
`(config-router)#address-family vpnv4 unicast`	Enter into vpnv4 unicast address family
`(config-router-af)#neighbor 2.2.2.2 activate`	Activate the neighbor inside vpnv4 address family
`(config-router-af)#exit-address-family`	Exit from address family vpnv4

© 2026 IP Infusion Inc. Proprietary

(config-router)# address-family ipv4 vrf vrf1	Enter into vrf address family
(config-router-af)# network 22.22.22.22/32	Advertise a route via iBGP connection.
(config-router-af)#redistribute ospf 1	Redistribute ospf routes into bgp
(config-router-af)#exit-address-family	Exit from address family vpnv4
(config-router)#exit	Exit from router BGP mode
(config)# commit	Committing the configuration to apply in running configuration
(config)#end	Exit from config mode into privilege mode

CE2 Configuration

#configure terminal	Enter the Configure mode.
(config)#interface lo	Enter interface mode
(config-if)#ip address 4.4.4.4/32 secondary	Configure the IP address of the interface loopback
(config-if)#exit	Exit interface mode
(config)# commit	Committing the configuration to apply in running configuration
(config)#interface xe5	Enter interface mode
(config-if)#ip address 20.20.40.39/31	Configure the IP address on the interface
(config-if)# description to_pe2	Adding Description to interface
(config-if)#exit	Exit interface mode.
(config)# commit	Committing the configuration to apply in running configuration
(config)#interface xe7	Enter interface mode
(config-if)#ip address 20.20.40.43/31	Configure the IP address on the interface
(config-if)# description ospf backdoor	Adding Description to interface
(config-if)# ip ospf cost 10	Set the OSPF cost of this link
(config-if)#exit	Exit interface mode.
(config)# commit	Committing the configuration to apply in running configuration
(config)#router ospf 1	Configure the routing process OSPF with process id
(config-router)#network 4.4.4.4/32 area 2	Define the interface on which OSPF runs, and associate the area ID with the interface.
(config-router)#network 20.20.40.38/31 area 2	Define the interface on which OSPF runs, and associate the area ID with the interface.
(config-router)#network 20.20.40.42/31 area 2	Define the interface on which OSPF runs, and associate the area ID with the interface.
(config-router)#exit	Exit from router ospf mode
(config)# commit	Committing the configuration to apply in running configuration

© 2026 IP Infusion Inc. Proprietary

Validation

CE1

```
CE1#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface           Instance ID
11.11.11.11       1   Full/Backup     00:00:34    20.20.40.41     xe4                     0
4.4.4.4           1   Full/Backup     00:00:36    20.20.40.43     xe24                    0

CE1#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, 00:08:20
O            4.4.4.4/32 [110/8] via 20.20.40.41, xe4, 00:01:48
O E2         11.11.11.11/32 [110/1] via 20.20.40.41, xe4, 00:01:48
O            20.20.40.38/31 [110/7] via 20.20.40.41, xe4, 00:01:48
C            20.20.40.40/31 is directly connected, xe4, 00:08:20
C            20.20.40.42/31 is directly connected, xe24, 00:08:20
O E2         22.22.22.22/32 [110/1] via 20.20.40.41, xe4, 00:02:22
C            127.0.0.0/8 is directly connected, lo, 00:13:50

Gateway of last resort is not set
```

PE1

```
PE1#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(vrf1):
Neighbor ID     Pri   State           Dead Time   Address         Interface           Instance ID
1.1.1.1           1   Full/DR         00:00:32    20.20.40.40     xe4                     0
22.22.22.22       1   Full/ -         00:00:35    22.22.22.22     SLINK0

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     00:00:34    20.20.40.49     xe10                    0
PE1#show ip ospf sham-links
Sham Link SLINK0 to destination 22.22.22.22 is up
  area 0.0.0.2 source 11.11.11.11
  Transmit Delay is 1 sec, State Point-To-Point,
  Timer intervals , Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:07
  No authentication
    Adjacency state Full
PE1#show ip ospf sham-links brief

  OSPF Process ID 1 VRF vrf1
  Total number of slinks: 1
  Shamlink Name    DestinationID  SourceID  Area    Status
  SLINK0          22.22.22.22          11.11.11.11 0.0.0.2  up
```

```
  OSPF Process ID 100 VRF default
  Total number of slinks: 0
  Shamlink Name    DestinationID  SourceID  Area    Status
PE1#show ip bgp vpnv4 all summary
BGP router identifier 2.2.2.2, 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
3.3.3.3                  4   100   16         17       2      0      0  00:03:49              4

Total number of neighbors 1

Total number of Established sessions 1
PE1#show ldp session
Peer IP Address           IF Name    My Role    State      KeepAlive  UpTime
3.3.3.3                   xe10       Passive   OPERATIONAL   30    00:03:39
PE1#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, 00:05:56
O IA         3.3.3.3/32 [110/2] via 20.20.40.49, xe10, 00:04:24
C            20.20.40.48/31 is directly connected, xe10, 00:05:56
C            127.0.0.0/8 is directly connected, lo, 00:16:19

Gateway of last resort is not set
```

PE2

```
PE2#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(vrf1):
Neighbor ID     Pri   State           Dead Time   Address         Interface            Instance ID
4.4.4.4           1   Full/DR         00:00:31    20.20.40.39     xe5                      0
11.11.11.11       1   Full/ -         00:00:36    11.11.11.11     SLINK0

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         00:00:39    20.20.40.48     xe10                     0
PE2#show ip bgp vpnv4 all summary
BGP router identifier 3.3.3.3, 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
2.2.2.2                  4   100   22         22       2      0      0  00:06:23              4

Total number of neighbors 1

Total number of Established sessions 1
PE2#show ip ospf sham-links
Sham Link SLINK0 to destination 11.11.11.11 is up
  area 0.0.0.2 source 22.22.22.22
```

```
  Transmit Delay is 1 sec, State Point-To-Point,
  Timer intervals , Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:03
  No authentication
    Adjacency state Full
PE2#show ip ospf sham-links brief

  OSPF Process ID 1 VRF vrf1
  Total number of slinks: 1
  Shamlink Name    DestinationID  SourceID  Area    Status
  SLINK0          11.11.11.11          22.22.22.22 0.0.0.2  up

  OSPF Process ID 100 VRF default
  Total number of slinks: 0
  Shamlink Name    DestinationID  SourceID  Area    Status
PE2#show ip bgp vpnv4 all summary
BGP router identifier 3.3.3.3, 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
2.2.2.2                  4   100   24         24       2      0      0  00:07:08               4

Total number of neighbors 1

Total number of Established sessions 1
PE2#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         2.2.2.2/32 [110/2] via 20.20.40.48, xe10, 00:07:32
C            3.3.3.3/32 is directly connected, lo, 00:07:52
C            20.20.40.48/31 is directly connected, xe10, 00:07:52
C            127.0.0.0/8 is directly connected, lo, 00:18:22

Gateway of last resort is not set
```

CE2

```
CE2#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface           Instance ID
22.22.22.22       1   Full/Backup     00:00:38    20.20.40.38     xe5                     0
1.1.1.1           1   Full/DR         00:00:35    20.20.40.42     xe7                     0
CE2#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            1.1.1.1/32 [110/8] via 20.20.40.38, xe5, 00:09:00
C            4.4.4.4/32 is directly connected, lo, 00:13:45
O E2         11.11.11.11/32 [110/1] via 20.20.40.38, xe5, 00:09:34
C            20.20.40.38/31 is directly connected, xe5, 00:13:45
O            20.20.40.40/31 [110/7] via 20.20.40.38, xe5, 00:09:00
C            20.20.40.42/31 is directly connected, xe7, 00:13:45
O E2         22.22.22.22/32 [110/1] via 20.20.40.38, xe5, 00:09:00
C            127.0.0.0/8 is directly connected, lo, 00:20:15

Gateway of last resort is not set
```

OSPF TE-Metric Extension

This section contains basic OSPF TE-Metric Extension configuration examples.

A TE link represents an IS-IS/OSPF link state advertisement and a link state database of certain physical resources and their properties between two nodes. Typically, a TE link is advertised as an adjunct to a "regular" OSPF or IS-IS link. That is, an adjacency is brought up on the link. When the link is up, both the regular IGP properties of the link (for example, the SPF metric) and the TE properties of the link are then advertised.

Note: Bandwidth-measurement values should be less than maximum link bandwidth configuration.

Topology

Figure 180. Basic OSPF 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 10.10.10.0/24 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 30.30.30.0/24 area 0`	Define the interface (30.30.30.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#commit`	Commit the configurations
`(config-router)#exit`	Exit mode
`(config)#interface eth1`	Interface mode
`(config-if)# delay-measurement uni-link-delay static 1`	Configure the delay-measurement uni-link-delay
`(config-if)# delay-measurement uni-link-delay a-bit-threshold min 2 max 4`	Configure the delay-measurement unilink delay a bit threshold
`(config-if)# bandwidth-measurement uni-available-bandwidth static 200k`	Configure the bandwidth-measurement unilink available bandwidth
`(config-if)#commit`	Commit the configurations
`(config-if)#exit`	Exit mode

© 2026 IP Infusion Inc. Proprietary

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`	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 20.20.20.0/24 area 0`	Define the interface (20.20.20.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).
`(config-router)#commit`	Commit the configurations
`(config-router)#exit`	Exit mode
`(config-if)# delay-measurement uni-link-delay static 1`	Configure the delay-measurement uni-link-delay
`(config-if)# delay-measurement uni-link-delay a-bit-threshold min 2 max 4`	Configure the delay-measurement unilink delay a bit threshold
`(config-if)# bandwidth-measurement uni-available-bandwidth static 200k`	Configure the bandwidth-measurement unilink available bandwidth
`(config-if)#commit`	Commit the configurations

Validation

```
 R1#show ip ospf database opaque-area

             OSPF Router with ID (20.20.20.1) (Process ID 100 VRF default)

                 Area-Local Opaque-LSA (Area 0.0.0.0)

   LS age: 1591
   Options: 0x22 (-|-|DC|-|-|-|E|-)
   LS Type: Area-Local Opaque-LSA
   Link State ID: 1.0.0.1 (Area-Local Opaque-Type/ID)
   Opaque Type: 1
   Opaque ID: 1
   Advertising Router: 20.20.20.1
   LS Seq Number: 80000004
   Checksum: 0x2a87
   Length: 28

     MPLS TE router ID : 20.20.20.1

     Number of Links : 0

   LS age: 1678
   Options: 0x22 (-|-|DC|-|-|-|E|-)
   LS Type: Area-Local Opaque-LSA
   Link State ID: 1.0.0.1 (Area-Local Opaque-Type/ID)
   Opaque Type: 1
   Opaque ID: 1
   Advertising Router: 30.30.30.1
   LS Seq Number: 80000004
   Checksum: 0x2a4b
```

```
  Length: 28

    MPLS TE router ID : 30.30.30.1

    Number of Links : 0

  LS age: 121
  Options: 0x22 (-|-|DC|-|-|-|E|-)
  LS Type: Area-Local Opaque-LSA
  Link State ID: 1.0.0.30 (Area-Local Opaque-Type/ID)
  Opaque Type: 1
  Opaque ID: 30
  Advertising Router: 20.20.20.1
  LS Seq Number: 80000006
  Checksum: 0xeabc
  Length: 116

    Link connected to Broadcast network
      Link ID : 10.10.10.2
      Interface Address : 10.10.10.1
      Admin Metric : 1
      Maximum bandwidth : 10000000.00 Kbits/s
      Maximum reservable bandwidth : 10000000.00 Kbits/s
      Unreserved Bandwidth :
      Number of Priority : 8
      Priority 0 : 10000000.00 Kbits/s       Priority 1 : 10000000.00 Kbits/s
      Priority 2 : 10000000.00 Kbits/s       Priority 3 : 10000000.00 Kbits/s
      Priority 4 : 10000000.00 Kbits/s       Priority 5 : 10000000.00 Kbits/s
      Priority 6 : 10000000.00 Kbits/s       Priority 7 : 10000000.00 Kbits/s
      Link Delay : 1 us, Anomalous : 0

    Number of Links : 1

  LS age: 36
  Options: 0x22 (-|-|DC|-|-|-|E|-)
  LS Type: Area-Local Opaque-LSA
  Link State ID: 1.0.0.32 (Area-Local Opaque-Type/ID)
  Opaque Type: 1
  Opaque ID: 32
  Advertising Router: 30.30.30.1
  LS Seq Number: 80000005
  Checksum: 0x0681
  Length: 116

    Link connected to Broadcast network
      Link ID : 10.10.10.2
      Interface Address : 10.10.10.2
      Admin Metric : 1
      Maximum bandwidth : 10000000.00 Kbits/s
      Maximum reservable bandwidth : 10000000.00 Kbits/s
      Unreserved Bandwidth :
      Number of Priority : 8
      Priority 0 : 10000000.00 Kbits/s       Priority 1 : 10000000.00 Kbits/s
      Priority 2 : 10000000.00 Kbits/s       Priority 3 : 10000000.00 Kbits/s
      Priority 4 : 10000000.00 Kbits/s       Priority 5 : 10000000.00 Kbits/s
      Priority 6 : 10000000.00 Kbits/s       Priority 7 : 10000000.00 Kbits/s
      Link Delay : 1 us, Anomalous : 0

    Number of Links : 1
R2    IP-Extended 21.21.21.0/24
R2# R2#sh ip ospf database

            OSPF Router with ID (30.30.30.1) (Process ID 100 VRF default)

                Router Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#       CkSum  Link count
20.20.20.1      20.20.20.1      1773     0x80000005 0xdd4a 2
```

```
30.30.30.1      30.30.30.1      1773     0x80000006 0xe9e1 2

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router       Age     Seq#       CkSum
10.10.10.2      30.30.30.1       133     0x80000002 0x0520

                Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID         ADV Router      Age      Seq#       CkSum  Opaque ID
1.0.0.1         20.20.20.1       138     0x80000002 0x2e85 1
1.0.0.1         30.30.30.1       173     0x80000002 0x2e49 1
1.0.0.30        20.20.20.1       362     0x80000002 0x5555 30
1.0.0.32        30.30.30.1       223     0x80000002 0x3288 32
R2#
R1#sh ip ospf database opaque-area

            OSPF Router with ID (20.20.20.1) (Process ID 100 VRF default)

                Area-Local Opaque-LSA (Area 0.0.0.0)

  LS age: 341
  Options: 0x22 (-|-|DC|-|-|-|E|-)
  LS Type: Area-Local Opaque-LSA
  Link State ID: 1.0.0.1 (Area-Local Opaque-Type/ID)
  Opaque Type: 1
  Opaque ID: 1
  Advertising Router: 20.20.20.1
  LS Seq Number: 80000006
  Checksum: 0x2689
  Length: 28

    MPLS TE router ID : 20.20.20.1

    Number of Links : 0

  LS age: 379
  Options: 0x22 (-|-|DC|-|-|-|E|-)
  LS Type: Area-Local Opaque-LSA
  Link State ID: 1.0.0.1 (Area-Local Opaque-Type/ID)
  Opaque Type: 1
  Opaque ID: 1
  Advertising Router: 30.30.30.1
  LS Seq Number: 80000006
  Checksum: 0x264d
  Length: 28

    MPLS TE router ID : 30.30.30.1

    Number of Links : 0

  LS age: 8
  Options: 0x22 (-|-|DC|-|-|-|E|-)
  LS Type: Area-Local Opaque-LSA
  Link State ID: 1.0.0.30 (Area-Local Opaque-Type/ID)
  Opaque Type: 1
  Opaque ID: 30
  Advertising Router: 20.20.20.1
  LS Seq Number: 8000000a
  Checksum: 0x1c03
  Length: 116

    Link connected to Broadcast network
      Link ID : 10.10.10.2
      Interface Address : 10.10.10.1
      Admin Metric : 1
      Maximum bandwidth : 10000000.00 Kbits/s
      Maximum reservable bandwidth : 10000000.00 Kbits/s
```

© 2026 IP Infusion Inc. Proprietary

```
      Unreserved Bandwidth :
      Number of Priority : 8
      Priority 0 : 10000000.00 Kbits/s       Priority 1 : 10000000.00 Kbits/s
      Priority 2 : 10000000.00 Kbits/s       Priority 3 : 10000000.00 Kbits/s
      Priority 4 : 10000000.00 Kbits/s       Priority 5 : 10000000.00 Kbits/s
      Priority 6 : 10000000.00 Kbits/s       Priority 7 : 10000000.00 Kbits/s
      Link Delay : 5 us, Anomalous : 1

    Number of Links : 1

  LS age: 33
  Options: 0x22 (-|-|DC|-|-|-|E|-)
  LS Type: Area-Local Opaque-LSA
  Link State ID: 1.0.0.32 (Area-Local Opaque-Type/ID)
  Opaque Type: 1
  Opaque ID: 32
  Advertising Router: 30.30.30.1
  LS Seq Number: 80000009
  Checksum: 0x37c7
  Length: 116

    Link connected to Broadcast network
      Link ID : 10.10.10.2
      Interface Address : 10.10.10.2
      Admin Metric : 1
      Maximum bandwidth : 10000000.00 Kbits/s
      Maximum reservable bandwidth : 10000000.00 Kbits/s
      Unreserved Bandwidth :
      Number of Priority : 8
      Priority 0 : 10000000.00 Kbits/s       Priority 1 : 10000000.00 Kbits/s
      Priority 2 : 10000000.00 Kbits/s       Priority 3 : 10000000.00 Kbits/s
      Priority 4 : 10000000.00 Kbits/s       Priority 5 : 10000000.00 Kbits/s
      Priority 6 : 10000000.00 Kbits/s       Priority 7 : 10000000.00 Kbits/s
      Link Delay : 5 us, Anomalous : 1

    Number of Links : 1
R1#sh ip ospf database opaque-area

            OSPF Router with ID (20.20.20.1) (Process ID 100 VRF default)

                Area-Local Opaque-LSA (Area 0.0.0.0)

  LS age: 341
  Options: 0x22 (-|-|DC|-|-|-|E|-)
  LS Type: Area-Local Opaque-LSA
  Link State ID: 1.0.0.1 (Area-Local Opaque-Type/ID)
  Opaque Type: 1
  Opaque ID: 1
  Advertising Router: 20.20.20.1
  LS Seq Number: 80000006
  Checksum: 0x2689
  Length: 28

    MPLS TE router ID : 20.20.20.1

    Number of Links : 0

  LS age: 379
  Options: 0x22 (-|-|DC|-|-|-|E|-)
  LS Type: Area-Local Opaque-LSA
  Link State ID: 1.0.0.1 (Area-Local Opaque-Type/ID)
  Opaque Type: 1
  Opaque ID: 1
  Advertising Router: 30.30.30.1
  LS Seq Number: 80000006
  Checksum: 0x264d
  Length: 28
```

```
    MPLS TE router ID : 30.30.30.1

    Number of Links : 0

  LS age: 8
  Options: 0x22 (-|-|DC|-|-|-|E|-)
  LS Type: Area-Local Opaque-LSA
  Link State ID: 1.0.0.30 (Area-Local Opaque-Type/ID)
  Opaque Type: 1
  Opaque ID: 30
  Advertising Router: 20.20.20.1
  LS Seq Number: 8000000a
  Checksum: 0x1c03
  Length: 116

    Link connected to Broadcast network
      Link ID : 10.10.10.2
      Interface Address : 10.10.10.1
      Admin Metric : 1
      Maximum bandwidth : 10000000.00 Kbits/s
      Maximum reservable bandwidth : 10000000.00 Kbits/s
      Unreserved Bandwidth :
      Number of Priority : 8
      Priority 0 : 10000000.00 Kbits/s       Priority 1 : 10000000.00 Kbits/s
      Priority 2 : 10000000.00 Kbits/s       Priority 3 : 10000000.00 Kbits/s
      Priority 4 : 10000000.00 Kbits/s       Priority 5 : 10000000.00 Kbits/s
      Priority 6 : 10000000.00 Kbits/s       Priority 7 : 10000000.00 Kbits/s
      Link Delay : 5 us, Anomalous : 1

    Number of Links : 1

  LS age: 33
  Options: 0x22 (-|-|DC|-|-|-|E|-)
  LS Type: Area-Local Opaque-LSA
  Link State ID: 1.0.0.32 (Area-Local Opaque-Type/ID)
  Opaque Type: 1
  Opaque ID: 32
  Advertising Router: 30.30.30.1
  LS Seq Number: 80000009
  Checksum: 0x37c7
  Length: 116

    Link connected to Broadcast network
      Link ID : 10.10.10.2
      Interface Address : 10.10.10.2
      Admin Metric : 1
      Maximum bandwidth : 10000000.00 Kbits/s
      Maximum reservable bandwidth : 10000000.00 Kbits/s
      Unreserved Bandwidth :
      Number of Priority : 8
      Priority 0 : 10000000.00 Kbits/s       Priority 1 : 10000000.00 Kbits/s
      Priority 2 : 10000000.00 Kbits/s       Priority 3 : 10000000.00 Kbits/s
      Priority 4 : 10000000.00 Kbits/s       Priority 5 : 10000000.00 Kbits/s
      Priority 6 : 10000000.00 Kbits/s       Priority 7 : 10000000.00 Kbits/s
      Link Delay : 5 us, Anomalous : 1

    Number of Links : 1
```

© 2026 IP Infusion Inc. Proprietary

OSPFv2 LDP Remote Loop-Free Alternate (RLFA)

A basic mechanism using Loop-Free Alternates (LFAs) is described in RFC5286 that provides good repair coverage in many topologies, especially those that are highly meshed.

However, some topologies, notably ring-based topologies, are not well protected by LFAs alone. This is because there is no neighbor of the Point of Local Repair (PLR) that has a cost to the destination via a path that does not traverse the failure that is cheaper than the cost to the destination via the failure.

RFC 7490 provides extensions to the basic repair mechanism in which tunnels are used to provide additional logical links that can be used as loop-free alternates where none exist in the original topology. It provides loop-free alternates that guarantee only link protection.

RFC 8102 provides remote-loop-free-based IP fast reroute mechanisms that specifies procedures for determining whether or not a given PQ-node provides node protection for a specific destination. It provides node protection for all destinations covered by the same remote-LFA alternate, in case of failure of the primary next-hop node

Ospf shall compute PQ node and LDP shall dynamically create tunnel to PQ node so that if primary path fails traffic can be rerouted to backup rLFA tunnel and hence to destination.

Remote LFA involves the use of a tunnel to a next-hop that is not directly connected. This is the primary difference between the LFA and Remote LFA.

RLFA Terminology

Note: Terms used for defining tunnel as below:

Figure 181. Understanding RLFA

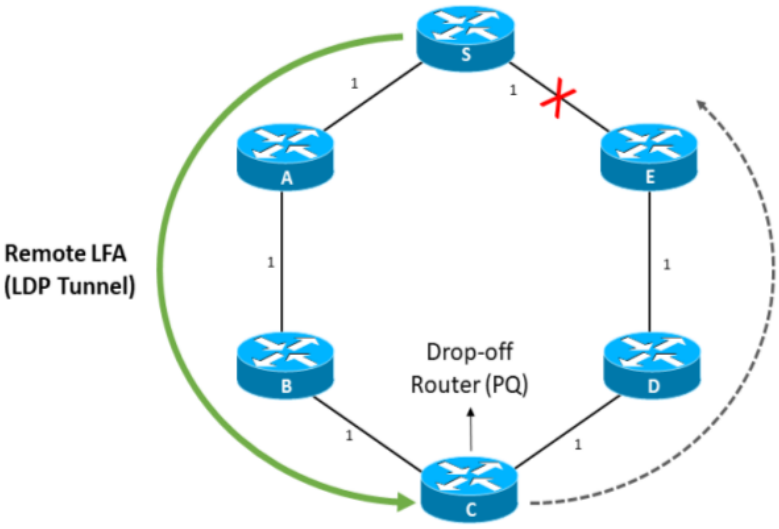

Repair tunnel

A tunnel established for the purpose of providing a virtual neighbor that is a Loop-Free Alternate.

P-space

The P-space of a router with respect to a protected link is the set of routers reachable from that specific router using the pre-convergence shortest paths, without any of those paths (including equal cost path splits) transiting that protected link.

For example, the P-space of S with respect to link S-E is the set of routers that S can reach without using the protected link S-E.

Extended P-space

Consider the set of neighbors of a router protecting a link. Exclude the router reachable over the protected link from that set of routers. The extended P-space of the protecting router with respect to the protected link is the union of the P spaces of the neighbors in that set of neighbors with respect to the protected link.

Q-space

Q-space of a router with respect to a protected link is the set of routers from which that specific router that can be reached without any path (including ECMP Splits) transiting that protected link.

© 2026 IP Infusion Inc. Proprietary

PQ node

A PQ node of a node S with respect to a protected link S-E is a node that is a member of both the P-space (or the extended P-space) of S with respect to that protected link S-E and the Q-space of E with respect to that protected link S-E. A repair tunnel endpoint is chosen from the set of PQ-nodes.

Remote LFA (RLFA)

The use of a PQ node rather than a neighbor of the repairing node as the next hop in an LFA repair.

In Figure 181, S can reach A, B, and C without going via S-E; these form S's extended P-space with respect to S-E. The routers that can reach E without going through S-E will be in E's Q-space with respect to link S-E; these are D and C.B has equal-cost paths to E via B-A-S-E and B-C-D-E, and so the forwarder at S might choose to send a packet to E via link S-E. Hence, B is not in the Q-space of E with respect to link S-E. The single node in both S's extended P-space and E's Q-space is C; thus, node C is selected as the repair tunnel's endpoint. thus, node C is selected as the repair tunnel's endpoint.

Establishing RLFA Tunnel

To calculate the Remote LFA backup path and to determines the Remote LFA node the software requirement can be broadly classified as below:

OSPF

OSPF shall calculate Repair path that are P space (routers it can reach without traversing the protected link) and Q space (routers that can reach the protected destination without traversing the protected link). Hence routers that belong to both spaces called as PQ routers has to be calculated. It has to inform LDP about PQ node so that RLFA tunnel can be established.

LDP

LDP shall establish targeted session with PQ node and shall advertise label to peer node for the destination FEC. LDP shall send primary and backup path FTN/ILM add/delete info to NSM for further programing.

NSM

NSM shall maintain Primary and Backup path FTN/ILM and shall send information to HSL for data-plane programing.

HSL

HSL program primary and backup FTN/ILM entry in hardware.

Configure LDP Remote Loop-Free Alternate (RLFA)

The LDP RLFA configuration process can be divided into the following tasks:

1. Enable label-switching on the interface on NSM.
2. Establish ospf routing between the nodes (to distribute reachability information within the MPLS cloud)
3. Configure ospf RLFA on Source node
4. Enabling LDP on an interface in the LDP daemon
5. Configure LDP FRR with Auto-targeted-session (Allow creating TLDP session dynamically)

© 2026 IP Infusion Inc. Proprietary

6. Enable BFD interval globally and for all ISIS enabled interfaces

Notes:

- Faster convergence can be achieved with lower BFD interval enabled globally.
- Dynamically created RLFA T-LDP sessions will be removed only after disabling LDP auto-targeted-session CLI or LDP FRR.
- When targeted-peer ipv4 CLI is configured with auto-targeted-session CLI, T-LDP session created for targeted-peer only remove after disabling auto-targeted-session CLI.
- After Enabling Ospf RLFA, Both LFA and RLFA computation will be done and RLFA path will be preferred to provide node-protection.
- Better convergence can be achieved with LDP-IGP-SYNC enabled.

Assumptions and Limitations

- RLFA Backup path computation will be supported only via IGP as IS-IS/Ospf.
- Only LDP(MPLS) will be used as a tunnel mechanism to reach a Remote-LFA repair node.
- Only IPv4 protocol is supported.
- RFC 7916 [LFA-MANAGE] is not supported.
- ECMP will not be supported for RLFA next-hop.

Topology

Figure 182 shows the configuration required to enable the RLFA feature.

Figure 182. RLFA Topology

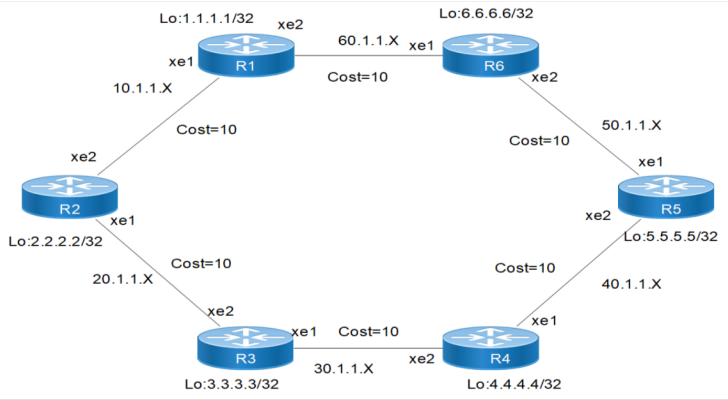

R1

`#configure terminal`	Enter configure mode
`(config)#router ldp`	Enable LDP process
`(config-router)#fast-reroute`	Enable LDP FRR

`(config-router)#auto-targeted-session`	To Allow creating TLDP session dynamically
`(config-router)#commit`	Commit the candidate configuration to the running configuration
`(config-router)#exit`	Exit LDP process
`(config)#router rsvp`	Enable RSPVP
`(config-router)#srlg-disjoint forced`	Configure SRLG to enable MPLS-TE
`(config-router)#commit`	Commit the candidate configuration to the running configuration
`(config)#interface xe1`	Enter interface mode.
`(config-if)#ip address 10.1.1.1/24`	Configure the IP address of the interface.
`(config-if)#ip ospf cost 10`	Configure the OSPF cost
`(config-if)#label-switching`	Enable label-switching on interface
`(config-if)#enable-ldp ipv4`	Enable LDP process on xe1 interface
`(config-if)#mpls traffic-eng srlg 11`	Enable MPLS-TE
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe2`	Enter interface mode.
`(config-if)#ip address 60.1.1.1/24`	Configure the IP address of the interface.
`(config-if)#ip ospf cost 10`	Configure OSPF cost
`(config-if)#label-switching`	Enable label-switching on interface
`(config-if)#mpls traffic-eng srlg 11`	Enable MPLS-TE
`(config-if)#enable-ldp ipv4`	Enable LDP process on xe2 interface
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode.
`(config)#interface lo`	Enter interface mode.
`(config-if)#ip address 1.1.1.1/32 secondary`	Configure the IP address of the interface
`(config-if)#exit`	Exit interface mode.
`(config)#router ospf 1`	Create an OSPF instance
`(config-router)# ospf router-id 1.1.1.1`	Configure router ID
`(config-router)# bfd all-interfaces`	Configure BFD
`(config-router)# timers spf exp 50 50`	Configure the OSPF timers
`(config-router)# timers throttle lsa all 0 1 1`	Configure the OSPF timer LSA throttle
`(config-router)# timers lsa arrival 1`	Configure the OSPF timer LSA arrival
`(config-router)# fast-reroute per-prefix remote-lfa area 0.0.0.0 tunnel mpls-ldp`	Configure Remote LFA to calculate backup paths to those destinations whichever does not satisfy basic LFA FRR inequalities

© 2026 IP Infusion Inc. Proprietary

(config-router)# network 1.1.1.1/32 area 0.0.0.0	Configure the network command to advertise the prefixes
(config-router)# network 10.1.1.0/24 area 0.0.0.0	Configure the network command to advertise the prefixes
(config-router)# network 60.1.1.0/24 area 0.0.0.0	Configure the network command to advertise the prefixes
(config-router)#commit	Commit the candidate configuration to the running configuration
(config-router)#end	Exit router mode.
(config)# bfd interval 3 minrx 3 multiplier 3	Configure BFD interval globally

R2

#configure terminal	Enter configure mode
(config)#router ldp	Enable LDP process
(config-router)#fast-reroute	Enable LDP FRR
(config-router)#auto-targeted-session	To Allow creating TLDP session dynamically
(config-router)#commit	Commit the candidate configuration to the running configuration
(config-router)#exit	Exit LDP process
(config)#router rsvp	Enable RSPVP
(config-router)#srlg-disjoint forced	Configure SRLG to enable mpls-TE
(config-router)#commit	Commit the candidate configuration to the running configuration
(config)#interface xe2	Enter interface mode.
(config-if)#ip address 10.1.1.2/24	Configure the IP address of the interface.
(config-if)#ip ospf cost 10	Configure the OSPF cost
(config-if)#label-switching	Enable label-switching on interface
(config-if)#enable-ldp ipv4	Enable LDP process on xe1 interface
(config-if)#mpls traffic-eng srlg 11	Enable MPLS-TE
(config-if)#commit	Commit the candidate configuration to the running configuration
(config-if)#exit	Exit interface mode.
(config)#interface xe1	Enter interface mode.
(config-if)#ip address 20.1.1.1/24	Configure the IP address of the interface.
(config-if)#ip ospf cost 10	Configure OSPF cost
(config-if)#label-switching	Enable label-switching on interface
(config-if)#mpls traffic-eng srlg 11	Enable MPLS-TE
(config-if)#enable-ldp ipv4	Enable LDP process on xe2 interface

© 2026 IP Infusion Inc. Proprietary

(config-if)#commit	Commit the candidate configuration to the running configuration
(config-if)#exit	Exit interface mode.
(config)#interface lo	Enter interface mode.
(config-if)#ip address 2.2.2.2/32 secondary	Configure the IP address of the interface
(config-if)#exit	Exit interface mode.
(config)#router ospf 1	Create an OSPF instance
(config-router)# ospf router-id 2.2.2.2	Configure router ID
(config-router)# bfd all-interfaces	Configure BFD
(config-router)# timers spf exp 50 50	Configure the OSPF timers
(config-router)# timers throttle lsa all 0 1 1	Configure the OSPF timer LSA throttle
(config-router)# timers lsa arrival 1	Configure the OSPF timer LSA arrival
(config-router)# network 2.2.2.2/32 area 0.0.0.0	Configure the network command to advertise the prefixes
(config-router)# network 10.1.1.0/24 area 0.0.0.0	Configure the network command to advertise the prefixes
(config-router)# network 20.1.1.0/24 area 0.0.0.0	Configure the network command to advertise the prefixes
(config-router)#commit	Commit the candidate configuration to the running configuration
(config)# bfd interval 3 minrx 3 multiplier 3 Configure bfd interval globally	Configure BFD interval globally
(config)#commit	Commit the candidate configuration to the running configuration
(config)#end	Exit configure mode.

R3

#configure terminal	Enter configure mode
(config)#router ldp	Enable LDP process
(config-router)#fast-reroute	Enable LDP FRR
(config-router)#auto-targeted-session	To Allow creating TLDP session dynamically
(config-router)#commit	Commit the candidate configuration to the running configuration
(config-router)#exit	Exit LDP process
(config)#router rsvp	Enable RSPVP
(config-router)#srlg-disjoint forced	Configure SRLG to enable mpls-TE
(config-router)#commit	Commit the candidate configuration to the running configuration

© 2026 IP Infusion Inc. Proprietary

`(config)#interface xe2`	Enter interface mode.
`(config-if)#ip address 20.1.1.2/24`	Configure the IP address of the interface.
`(config-if)#ip ospf cost 10`	Configure the OSPF cost
`(config-if)#label-switching`	Enable label-switching on interface
`(config-if)#enable-ldp ipv4`	Enable LDP process on xe1 interface
`(config-if)#mpls traffic-eng srlg 11`	Enable MPLS-TE
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe1`	Enter interface mode.
`(config-if)#ip address 30.1.1.1/24`	Configure the IP address of the interface.
`(config-if)#ip ospf cost 10`	Configure OSPF cost
`(config-if)#label-switching`	Enable label-switching on interface
`(config-if)#mpls traffic-eng srlg 11`	Enable MPLS-TE
`(config-if)#enable-ldp ipv4`	Enable ldp process on xe2 interface
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode.
`(config)#interface lo`	Enter interface mode.
`(config-if)#ip address 3.3.3.3/32 secondary`	Configure the IP address of the interface
`(config-if)#exitv`	Exit interface mode.
`(config)#router ospf 1`	Create an OSPF instance
`(config-router)# ospf router-id 3.3.3.3`	Configure router id
`(config-router)# bfd all-interfaces`	Configure BFD
`(config-router)# timers spf exp 50 50`	Configure the OSPF timers
`(config-router)# timers throttle lsa all 0 1 1`	Configure the OSPF timer LSA throttle
`(config-router)# timers lsa arrival 1`	Configure the OSPF timer LSA arrival
`(config-router)# network 3.3.3.3/32 area 0.0.0.0`	Configure the network command to advertise the prefixes
`(config-router)# network 20.1.1.0/24 area 0.0.0.0`	Configure the network command to advertise the prefixes
`(config-router)# network 30.1.1.0/24 area 0.0.0.0`	Configure the network command to advertise the prefixes
`(config-router)#commit`	Commit the candidate configuration to the running configuration
`(config)# bfd interval 3 minrx 3 multiplier 3` `Configure bfd interval globally`	Configure BFD interval globally
`(config)#commit`	Commit the candidate configuration to the running

	configuration
(config)#end	Exit configure mode.

R4

#configure terminal	Enter configure mode
(config)#router ldp	Enable LDP process
(config-router)#fast-reroute	Enable LDP FRR
(config-router)#auto-targeted-session	To Allow creating TLDP session dynamically
(config-router)#commit	Commit the candidate configuration to the running configuration
(config-router)#exit	Exit LDP process
(config)#router rsvp	Enable RSPVP
(config-router)#srlg-disjoint forced	Configure SRLG to enable MPLS-TE
(config-router)#commit	Commit the candidate configuration to the running configuration
(config)#interface xe2	Enter interface mode.
(config-if)#ip address 30.1.1.2/24	Configure the IP address of the interface.
(config-if)#ip ospf cost 10	Configure the OSPF cost
(config-if)#label-switching	Enable label-switching on interface
(config-if)#enable-ldp ipv4	Enable LDP process on xe1 interface
(config-if)#mpls traffic-eng srlg 11	Enable MPLS-TE
(config-if)#commit	Commit the candidate configuration to the running configuration
(config-if)#exit	Exit interface mode.
(config)#interface xe1	Enter interface mode.
(config-if)#ip address 40.1.1.1/24	Configure the IP address of the interface.
(config-if)#ip ospf cost 10	Configure OSPF cost
(config-if)#label-switching	Enable label-switching on interface
(config-if)#mpls traffic-eng srlg 11	Enable MPLS-TE
(config-if)#enable-ldp ipv4	Enable LDP process on xe2 interface
(config-if)#commit	Commit the candidate configuration to the running configuration
(config-if)#exit	Exit interface mode.
(config)#interface lo	Enter interface mode.
(config-if)#ip address 4.4.4.4/32 secondary	Configure the IP address of the interface
(config-if)#exit	Exit interface mode.
(config)#router ospf 1	Create an OSPF instance

© 2026 IP Infusion Inc. Proprietary

(config-router)# ospf router-id 4.4.4.4	Configure router ID
(config-router)# bfd all-interfaces	Configure BFD
(config-router)# timers spf exp 50 50	Configure the OSPF timers
(config-router)# timers throttle lsa all 0 1 1	Configure the OSPF timer LSA throttle
(config-router)# timers lsa arrival 1	Configure the OSPF timer LSA arrival
(config-router)# network 4.4.4.4/32 area 0.0.0.0	Configure the network command to advertise the prefixes
(config-router)# network 40.1.1.0/24 area 0.0.0.0	Configure the network command to advertise the prefixes
(config-router)# network 30.1.1.0/24 area 0.0.0.0	Configure the network command to advertise the prefixes
(config-router)#commit	Commit the candidate configuration to the running configuration
(config)# bfd interval 3 minrx 3 multiplier 3 Configure bfd interval globally	Configure BFD interval globally
(config)#commit	Commit the candidate configuration to the running configuration
(config)#end	Exit configure mode.

R5

#configure terminal	Enter configure mode
(config)#router ldp	Enable LDP process
(config-router)#fast-reroute	Enable LDP FRR
(config-router)#auto-targeted-session	To Allow creating TLDP session dynamically
(config-router)#commit	Commit the candidate configuration to the running configuration
(config-router)#exit	Exit LDP process
(config)#router rsvp	Enable RSPVP
(config-router)#srlg-disjoint forced	Configure SRLG to enable mpls-TE
(config-router)#commit	Commit the candidate configuration to the running configuration
(config)#interface xe2	Enter interface mode.
(config-if)#ip address 40.1.1.2/24	Configure the IP address of the interface.
(config-if)#ip ospf cost 10	Configure the OSPF cost
(config-if)#label-switching	Enable label-switching on interface
(config-if)#enable-ldp ipv4	Enable LDP process on xe1 interface
(config-if)#mpls traffic-eng srlg 11	Enable MPLS-TE
(config-if)#commit	Commit the candidate configuration to the running configuration

© 2026 IP Infusion Inc. Proprietary

(config-if)#exit	Exit interface mode.
(config)#interface xe1	Enter interface mode.
(config-if)#ip address 50.1.1.1/24	Configure the IP address of the interface.
(config-if)#ip ospf cost 10	Configure OSPF cost
(config-if)#label-switching	Enable label-switching on interface
(config-if)#mpls traffic-eng srlg 11	Enable MPLS-TE
(config-if)#enable-ldp ipv4	Enable LDP process on xe2 interface
(config-if)#commit	Commit the candidate configuration to the running configuration
(config-if)#exit	Exit interface mode.
(config)#interface lo	Enter interface mode.
(config-if)#ip address 5.5.5.5/32 secondary	Configure the IP address of the interface
(config-if)#exit	Exit interface mode.
(config)#router ospf 1	Create an OSPF instance
(config-router)# ospf router-id 5.5.5.5	Configure router ID
(config-router)# bfd all-interfaces	Configure BFD
(config-router)# timers spf exp 50 50	Configure the OSPF timers
(config-router)# timers throttle lsa all 0 1 1	Configure the OSPF timer LSA throttle
(config-router)# timers lsa arrival 1	Configure the OSPF timer LSA arrival
(config-router)# network 5.5.5.5/32 area 0.0.0.0	Configure the network command to advertise the prefixes
(config-router)# network 50.1.1.0/24 area 0.0.0.0	Configure the network command to advertise the prefixes
(config-router)# network 40.1.1.0/24 area 0.0.0.0	Configure the network command to advertise the prefixes
(config-router)#commit	Commit the candidate configuration to the running configuration
(config)# bfd interval 3 minrx 3 multiplier 3 Configure bfd interval globally	Configure BFD interval globally
(config)#commit	Commit the candidate configuration to the running configuration
(config)#end	Exit configure mode.

R6

#configure terminal	Enter configure mode
(config)#router ldp	Enable LDP process
(config-router)#fast-reroute	Enable LDP FRR
(config-router)#auto-targeted-session	To Allow creating TLDP session dynamically
(config-router)#commit	Commit the candidate configuration to the running

© 2026 IP Infusion Inc. Proprietary

	configuration
(config-router)#exit	Exit LDP process
(config)#router rsvp	Enable RSPVP
(config-router)#srlg-disjoint forced	Configure SRLG to enable mpls-TE
(config-router)#commit	Commit the candidate configuration to the running configuration
(config)#interface xe2	Enter interface mode.
(config-if)#ip address 50.1.1.2/24	Configure the IP address of the interface.
(config-if)#ip ospf cost 10	Configure the OSPF cost
(config-if)#label-switching	Enable label-switching on interface
(config-if)#enable-ldp ipv4	Enable LDP process on xe1 interface
(config-if)#mpls traffic-eng srlg 11	Enable MPLS-TE
(config-if)#commit	Commit the candidate configuration to the running configuration
(config-if)#exit	Exit interface mode.
(config)#interface xe1	Enter interface mode.
(config-if)#ip address 60.1.1.2/24	Configure the IP address of the interface.
(config-if)#ip ospf cost 10	Configure OSPF cost
(config-if)#label-switching	Enable label-switching on interface
(config-if)#mpls traffic-eng srlg 11	Enable MPLS-TE
(config-if)#enable-ldp ipv4	Enable LDP process on xe2 interface
(config-if)#commit	Commit the candidate configuration to the running configuration
(config-if)#exit	Exit interface mode.
(config)#interface lo	Enter interface mode.
(config-if)#ip address 6.6.6.6/32 secondary	Configure the IP address of the interface
(config-if)#exit	Exit interface mode.
(config)#router ospf 1	Create an OSPF instance
(config-router)# ospf router-id 6.6.6.6	Configure router id
(config-router)# bfd all-interfaces	Configure BFD
(config-router)# timers spf exp 50 50	Configure the OSPF timers
(config-router)# timers throttle lsa all 0 1 1	Configure the OSPF timer LSA throttle
(config-router)# timers lsa arrival 1	Configure the OSPF timer LSA arrival
(config-router)# network 6.6.6.6/32 area 0.0.0.0	Configure the network command to advertise the prefixes
(config-router)# network 50.1.1.0/24 area 0.0.0.0	Configure the network command to advertise the prefixes

© 2026 IP Infusion Inc. Proprietary

`(config-router)# network 60.1.1.0/24 area 0.0.0.0`	Configure the network command to advertise the prefixes
`(config-router)#commit`	Commit the candidate configuration to the running configuration
`(config)# bfd interval 3 minrx 3 multiplier 3` `Configure bfd interval globally`	Configure BFD interval globally
`(config)#commit`	Commit the candidate configuration to the running configuration
`(config)#end`	Exit configure mode.

Validation

R1

Check LDP neighborship before enabling RLFA.

```
#show ldp session
Peer IP Address          IF Name    My Role    State      KeepAlive  UpTime
2.2.2.2                  xe20       Passive   OPERATIONAL   30    00:02:43
6.6.6.6                  ge10       Passive   OPERATIONAL   30    00:00:30
```

Check the output of `show ip ospf neighbors` to verify that OSPF adjacency is up.

```
#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID     Pri   State            Dead Time   Address         Interface
       Instance ID
2.2.2.2           1   Full/Backup      00:00:39    10.1.1.2        xe20
           0
6.6.6.6           1   Full/DR          00:00:35    60.1.1.2        ge10
           0
```

Check the OSPF route installation in the OSPF table and RIB table.

```
#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, 00:37:05
O            2.2.2.2/32 [110/11] via 10.1.1.2, xe20, 00:04:49
O            3.3.3.3/32 [110/21] via 10.1.1.2, xe20, 00:04:49
O            4.4.4.4/32 [110/31] via 60.1.1.2, ge10, 00:04:49
                        [110/31] via 10.1.1.2, xe20
O            5.5.5.5/32 [110/21] via 60.1.1.2, ge10, 00:02:29
O            6.6.6.6/32 [110/11] via 60.1.1.2, ge10, 00:02:29
C            10.1.1.0/24 is directly connected, xe20, 00:33:59
O            20.1.1.0/24 [110/20] via 10.1.1.2, xe20, 00:04:49
O            30.1.1.0/24 [110/30] via 10.1.1.2, xe20, 00:04:49
O            40.1.1.0/24 [110/30] via 60.1.1.2, ge10, 00:02:29
O            50.1.1.0/24 [110/20] via 60.1.1.2, ge10, 00:02:29
C            60.1.1.0/24 is directly connected, ge10, 00:02:36
C            127.0.0.0/8 is directly connected, lo, 00:45:19
```

© 2026 IP Infusion Inc. Proprietary

```
Gateway of last resort is not set

#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
       OSPF LFA attributes:
       P - Primary, SP - Secondary-Path, LP - Link Protecting,
       NP - Node Protecting, BID - Broadcast Link Protecting
       DP - Downstream Protecting

C  1.1.1.1/32 [1] is directly connected, lo, Area 0.0.0.0
O  2.2.2.2/32 [11] via 10.1.1.2, xe20, Area 0.0.0.0
O  3.3.3.3/32 [21] via 10.1.1.2, xe20, Area 0.0.0.0
O  4.4.4.4/32 [31] via 10.1.1.2, xe20, Area 0.0.0.0
                   via 60.1.1.2, ge10, Area 0.0.0.0
O  5.5.5.5/32 [21] via 60.1.1.2, ge10, Area 0.0.0.0
O  6.6.6.6/32 [11] via 60.1.1.2, ge10, Area 0.0.0.0
C  10.1.1.0/24 [10] is directly connected, xe20, Area 0.0.0.0
O  20.1.1.0/24 [20] via 10.1.1.2, xe20, Area 0.0.0.0
O  30.1.1.0/24 [30] via 10.1.1.2, xe20, Area 0.0.0.0
O  40.1.1.0/24 [30] via 60.1.1.2, ge10, Area 0.0.0.0
O  50.1.1.0/24 [20] via 60.1.1.2, ge10, Area 0.0.0.0
C  60.1.1.0/24 [10] is directly connected, ge10, Area 0.0.0.0
```

Verify OSPF LFA and RLFA backup computed paths for Primary Paths

```
#show ip ospf route fast-reroute

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
       OSPF LFA attributes:
       P - Primary, SP - Secondary-Path, LP - Link Protecting,
       NP - Node Protecting, BID - Broadcast Link Protecting
       DP - Downstream Protecting

O  2.2.2.2/32 [11] via 10.1.1.2, xe20, Area 0.0.0.0
                       Remote FRR path:
                       via 4.4.4.4, via 60.1.1.2, ge10, Area 0.0.0.0
                      Attributes: Metric: [51] ,LP
O  3.3.3.3/32 [21] via 10.1.1.2, xe20, Area 0.0.0.0
                       Remote FRR path:
                       via 4.4.4.4, via 60.1.1.2, ge10, Area 0.0.0.0
                      Attributes: Metric: [41] ,LP ,NP ,DP
O  4.4.4.4/32 [31] via 10.1.1.2, xe20, Area 0.0.0.0
                       Backup path:
                        via 60.1.1.2, ge10, Area 0.0.0.0
                       Attributes: Metric: [31] ,P ,NP ,BID ,DP
                  via 60.1.1.2, ge10, Area 0.0.0.0
                       Backup path:
                        via 10.1.1.2, xe20, Area 0.0.0.0
                       Attributes: Metric: [31] ,P ,NP ,BID ,DP
O  5.5.5.5/32 [21] via 60.1.1.2, ge10, Area 0.0.0.0
                       Remote FRR path:
                       via 4.4.4.4, via 10.1.1.2, xe20, Area 0.0.0.0
                      Attributes: Metric: [41] ,LP ,NP ,DP
O  6.6.6.6/32 [11] via 60.1.1.2, ge10, Area 0.0.0.0
                       Remote FRR path:
                       via 4.4.4.4, via 10.1.1.2, xe20, Area 0.0.0.0
                      Attributes: Metric: [51] ,LP
O  20.1.1.0/24 [20] via 10.1.1.2, xe20, Area 0.0.0.0
                       Remote FRR path:
                        via 4.4.4.4, via 60.1.1.2, ge10, Area 0.0.0.0
```

```
                        Attributes: Metric: [50] ,LP ,NP
O  30.1.1.0/24 [30] via 10.1.1.2, xe20, Area 0.0.0.0
                        Backup path:
                          via 60.1.1.2, ge10, Area 0.0.0.0
                         Attributes: Metric: [40] ,SP ,NP ,BID
O  40.1.1.0/24 [30] via 60.1.1.2, ge10, Area 0.0.0.0
                        Backup path:
                          via 10.1.1.2, xe20, Area 0.0.0.0
                         Attributes: Metric: [40] ,SP ,NP ,BID
O  50.1.1.0/24 [20] via 60.1.1.2, ge10, Area 0.0.0.0
                        Remote FRR path:
                         via 4.4.4.4, via 10.1.1.2, xe20, Area 0.0.0.0
                        Attributes: Metric: [50] ,LP ,NP
```

Verify PQ node which is near to source is selected and Target-LDP session is established with PQ node using below commands

```
#show ldp session
Peer IP Address           IF Name    My Role    State      KeepAlive  UpTime
2.2.2.2                   xe20       Passive   OPERATIONAL   30    00:07:29
6.6.6.6                   ge10       Passive   OPERATIONAL   30    00:05:16
4.4.4.4                   xe20       Passive   OPERATIONAL   30    00:05:04
```

Verify that Primary and Backup FTN's are installed with labels in LDP RLFA route table

```
#sh ldp rlfa-routes
Codes: p - stale rLFA route
Fec             Primary-NH      Backup-NH       rLFA-Addr       Out-Intf  Outer-label Inner-label
Owner

2.2.2.2         10.1.1.2        60.1.1.2        4.4.4.4         ge10      24964       24969       osp
f
3.3.3.3         10.1.1.2        60.1.1.2        4.4.4.4         ge10      24964       24970       osp
f
5.5.5.5         60.1.1.2        10.1.1.2        4.4.4.4         xe20      24963       24971       osp
f
6.6.6.6         60.1.1.2        10.1.1.2        4.4.4.4         xe20      24963       24972       osp
f
20.1.1.0        10.1.1.2        60.1.1.2        4.4.4.4         ge10      24964       24974       osp
f
50.1.1.0        60.1.1.2        10.1.1.2        4.4.4.4         xe20      24963       24975       osp
f
```

Verify that backup XC's calculated for primary FTN's in MPLS forwarding table. Verify the same in FTN table.

```
#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
       B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport

Code    FEC                FTN-ID    Nhlfe-ID  Tunnel-id   Pri   LSP-Type     Out-Label    Out-
Intf    ELC   Nexthop
   L>   2.2.2.2/32          1         2         -           Yes   LSP_
DEFAULT  3            xe20        No    10.1.1.2
                                     23        -           No    LSP_
DEFAULT  24969        ge10        No    4.4.4.4                                                (via
60.1.1.2 ,label 24964)
   L>   3.3.3.3/32          2         4         -           Yes   LSP_
DEFAULT  24962        xe20        No    10.1.1.2
                                     24        -           No    LSP_
DEFAULT  24970        ge10        No    4.4.4.4                                                (via
60.1.1.2 ,label 24964)
   L>   4.4.4.4/32          3         6         -           Yes   LSP_
DEFAULT  24963        xe20        No    10.1.1.2
                                     7         -           No    LSP_
```

```
DEFAULT  24964        ge10        No    60.1.1.2
                                      7          -           Yes   LSP_
DEFAULT  24964        ge10        No    60.1.1.2
                                      5          -           No    LSP_
DEFAULT  24963        xe20        No    10.1.1.2
   L>   5.5.5.5/32          4          10         -           Yes   LSP_
DEFAULT  24965        ge10        No    60.1.1.2
                                      25         -           No    LSP_
DEFAULT  24971        xe20        No    4.4.4.4
                                                                                (via
10.1.1.2 ,label 24963)
   L>   6.6.6.6/32          5          15         -           Yes   LSP_
DEFAULT  3            ge10        No    60.1.1.2
                                      26         -           No    LSP_
DEFAULT  24972        xe20        No    4.4.4.4
                                                                                (via
10.1.1.2 ,label 24963)
   L>   20.1.1.0/24         6          11         -           Yes   LSP_
DEFAULT  3            xe20        No    10.1.1.2
                                      27         -           No    LSP_
DEFAULT  24974        ge10        No    4.4.4.4
                                                                                (via
60.1.1.2 ,label 24964)
   L>   30.1.1.0/24         7          13         -           Yes   LSP_
DEFAULT  24966        xe20        No    10.1.1.2
                                      16         -           No    LSP_
DEFAULT  24966        ge10        No    60.1.1.2
   L>   40.1.1.0/24         8          19         -           Yes   LSP_
DEFAULT  24967        ge10        No    60.1.1.2
                                      20         -           No    LSP_
DEFAULT  24967        xe20        No    10.1.1.2
   L>   50.1.1.0/24         9          22         -           Yes   LSP_
DEFAULT  3            ge10        No    60.1.1.2
                                      28         -           No    LSP_
DEFAULT  24975        xe20        No    4.4.4.4
                                                                                (via
10.1.1.2 ,label 24963)

#show mpls ftn-table
 Primary FTN entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
     Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
      Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 1, owner: N/A, Stale: NO, out intf: xe20, out label: 3
    Nexthop addr: 10.1.1.2        cross connect ix: 1, op code: Push

     Backup Cross connect ix: 9, in intf: - in label: 0  out-segment ix: 23
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 23, owner: LDP, Stale: NO, out intf: ge10, out label: 24969
    Nexthop addr: 4.4.4.4        cross connect ix: 9, op code: Push


 Primary FTN entry with FEC: 3.3.3.3/32, id: 2, row status: Active, Tunnel-Policy: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
     Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 3
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: xe20, out label: 24962
    Nexthop addr: 10.1.1.2        cross connect ix: 2, op code: Push

     Backup Cross connect ix: 11, in intf: - in label: 0  out-segment ix: 24
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 24, owner: LDP, Stale: NO, out intf: ge10, out label: 24970
    Nexthop addr: 4.4.4.4        cross connect ix: 11, op code: Push


 Primary FTN entry with FEC: 4.4.4.4/32, id: 3, row status: Active, Tunnel-Policy: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
```

```
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
    Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: xe20, out label: 24963
   Nexthop addr: 10.1.1.2        cross connect ix: 3, op code: Push

    Backup Cross connect ix: 2, in intf: - in label: 0  out-segment ix: 7
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: ge10, out label: 24964
   Nexthop addr: 60.1.1.2        cross connect ix: 3, op code: Push

    Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 7
     Owner: LDP, Persistent: No, Admin Status: Down, Oper Status: Not present
      Out-segment with ix: 7, owner: LDP, Stale: NO, out intf: ge10, out label: 24964
   Nexthop addr: 60.1.1.2        cross connect ix: 3, op code: Push

    Backup Cross connect ix: 1, in intf: - in label: 0  out-segment ix: 5
     Owner: LDP, Persistent: No, Admin Status: Down, Oper Status: Not present
      Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: xe20, out label: 24963
   Nexthop addr: 10.1.1.2        cross connect ix: 3, op code: Push


Primary FTN entry with FEC: 5.5.5.5/32, id: 4, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
    Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 9
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 9, owner: LDP, Stale: NO, out intf: ge10, out label: 24965
   Nexthop addr: 60.1.1.2        cross connect ix: 5, op code: Push

    Backup Cross connect ix: 13, in intf: - in label: 0  out-segment ix: 25
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 25, owner: LDP, Stale: NO, out intf: xe20, out label: 24971
   Nexthop addr: 4.4.4.4        cross connect ix: 13, op code: Push


Primary FTN entry with FEC: 6.6.6.6/32, id: 5, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
    Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 14
     Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 14, owner: N/A, Stale: NO, out intf: ge10, out label: 3
   Nexthop addr: 60.1.1.2        cross connect ix: 7, op code: Push

    Backup Cross connect ix: 15, in intf: - in label: 0  out-segment ix: 26
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 26, owner: LDP, Stale: NO, out intf: xe20, out label: 24972
   Nexthop addr: 4.4.4.4        cross connect ix: 15, op code: Push


Primary FTN entry with FEC: 20.1.1.0/24, id: 6, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
    Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
     Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: N/A, Stale: NO, out intf: xe20, out label: 3
   Nexthop addr: 10.1.1.2        cross connect ix: 1, op code: Push

    Backup Cross connect ix: 17, in intf: - in label: 0  out-segment ix: 27
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 27, owner: LDP, Stale: NO, out intf: ge10, out label: 24974
   Nexthop addr: 4.4.4.4        cross connect ix: 17, op code: Push


Primary FTN entry with FEC: 30.1.1.0/24, id: 7, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
    Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 12
```

© 2026 IP Infusion Inc. Proprietary

```
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 12, owner: LDP, Stale: NO, out intf: xe20, out label: 24966
   Nexthop addr: 10.1.1.2        cross connect ix: 6, op code: Push

    Backup Cross connect ix: 3, in intf: - in label: 0  out-segment ix: 16
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 16, owner: LDP, Stale: NO, out intf: ge10, out label: 24966
   Nexthop addr: 60.1.1.2        cross connect ix: 3, op code: Push


Primary FTN entry with FEC: 40.1.1.0/24, id: 8, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
    Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 18
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 18, owner: LDP, Stale: NO, out intf: ge10, out label: 24967
   Nexthop addr: 60.1.1.2        cross connect ix: 8, op code: Push

    Backup Cross connect ix: 4, in intf: - in label: 0  out-segment ix: 20
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 20, owner: LDP, Stale: NO, out intf: xe20, out label: 24967
   Nexthop addr: 10.1.1.2        cross connect ix: 4, op code: Push


Primary FTN entry with FEC: 50.1.1.0/24, id: 9, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
    Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 14
     Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 14, owner: N/A, Stale: NO, out intf: ge10, out label: 3
   Nexthop addr: 60.1.1.2        cross connect ix: 7, op code: Push

    Backup Cross connect ix: 19, in intf: - in label: 0  out-segment ix: 28
     Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 28, owner: LDP, Stale: NO, out intf: xe20, out label: 24975
   Nexthop addr: 4.4.4.4        cross connect ix: 19, op code: Push
```

OPEN SHORTEST PATH FIRST COMMAND REFERENCE

OSPFv2 Commands 2382
area authentication 2386
area default-cost 2387
area filter-list 2388
area interface 2390
area interface authentication 2391
area interface cost 2392
area interface database-filter 2393
area interface dead-interval 2394
area interface hello-interval 2395
area interface IFNAME bfd 2396
area interface IFNAME priority 2397
area interface mpls ldp-igp 2398
area interface mpls ldp-igp sync ospf holddown-timer 2399
area interface network-type 2400
area interface passive 2401
area interface retransmit-interval 2402
area interface te-metric 2403
area nssa 2404
area range 2406
area sham-link 2408
area shortcut 2410
area stub 2411
area virtual-link 2412
auto-cost reference bandwidth 2414
bfd all-interfaces 2415
capability cspf 2416
capability lls 2417
capability opaque 2418
capability traffic-engineering 2419
capability vrf-lite 2420
clear ip ospf 2421
compatible rfc1583 2422
debug ip ospf dist-ls 2423
debug ip ospf lfa 2424

© 2026 IP Infusion Inc. Proprietary

debug ip ospf redist 2425
debug ip ospf retransmission 2426
debug ospf 2427
debug ospf database-timer rate-limit 2429
debug ospf events 2430
debug ospf ifsm 2431
debug ospf lsa 2432
debug ospf nfsm 2433
debug ospf nsm 2434
debug ospf packet 2435
debug ospf rib 2436
debug ospf route 2437
default-information originate 2438
default-metric 2440
distance 2441
distribute bgp-ls 2443
distribute-list 2444
enable db-summary-opt 2446
enable ext-ospf-multi-inst 2447
fast-reroute keep-all-paths 2448
fast-reroute per-prefix remote-lfa area 2449
fast-reroute per-prefix route-map 2450
fast-reroute terminate-hold-on interval 2451
fast-reroute tie-break 2452
host area 2454
ip ospf authentication 2455
ip ospf authentication-key 2456
ip ospf bfd 2458
ip ospf cost 2459
ip ospf database-filter 2460
ip ospf dead-interval 2461
ip ospf demand-circuit 2462
ip ospf disable 2463
ip ospf fast-reroute per-prefix candidate disable 2464
ip ospf flood-reduction 2465
ip ospf hello-interval 2466
ip ospf message-digest-key 2467
ip ospf mtu 2469
ip ospf mtu-ignore 2470
ip ospf multi-area 2471

© 2026 IP Infusion Inc. Proprietary

ip ospf network 2473
ip ospf priority 2474
ip ospf resync-timeout 2475
ip ospf retransmit-interval 2476
ip ospf transmit-delay 2477
log-adjacency-changes 2478
max-concurrent-dd 2479
maximum-area 2480
max-metric 2481
neighbor 2483
network 2484
ospf abr-type 2486
ospf area-interface-config-mode 2487
ospf flood-reduction 2488
ospf point-point rfc-incompatible 2489
ospf router-id 2490
overflow database 2491
overflow database external 2492
passive-interface 2493
redistribute 2495
router ospf 2497
show cspf lsp 2497
show cspf rsvp forwarding-timer 2500
show debugging ospf 2501
show ip ospf 2502
show ip ospf border-routers 2506
show ip ospf database brief 2508
show ip ospf database detail 2510
show ip ospf igp-shortcut-lsp 2518
show ip ospf igp-shortcut-route 2519
show ip ospf interface 2520
show ip ospf multi-area-adjacencies 2523
show ip ospf neighbor 2525
show ip ospf route 2529
show ip ospf sham-links 2532
show ip ospf valid 2534
show ip ospf virtual-links 2535
show ip protocols 2537
show ip route fast-reroute 2539
shutdown 2540

© 2026 IP Infusion Inc. Proprietary

snmp context-name 2541
snmp restart ospf 2542
summary-address 2543
te-metric 2545
timers lsa arrival 2546
timers spf exp 2547
timers throttle lsa 2548
OSPFv2 Graceful Restart Commands 2549
action interface 2550
capability restart 2551
event nbr 2552
debug ip ospf graceful-restart 2553
if-event-track ospf event 2554
ospf restart grace-period 2555
ospf restart helper 2556
restart ip ospf graceful 2557
OSPFv3 Graceful Restart Commands 2558
capability restart 2559
ipv6 ospf restart grace-period 2560
ipv6 ospf restart helper 2561
ipv6 ospf restart planned-only 2562
restart ipv6 ospf graceful 2563
OSPFv3 Commands 2564
abr-type 2566
area default-cost 2567
area nssa 2568
area range 2570
area stub 2571
area virtual-link 2572
auto-cost reference bandwidth 2574
bfd all-interfaces 2575
clear ipv6 ospf process 2576
debug ipv6 ospf 2577
debug ipv6 ospf bfd 2578
debug ipv6 ospf events 2579
debug ipv6 ospf ifsm 2580
debug ipv6 ospf lsa 2581
debug ipv6 ospf nfsm 2582
debug ipv6 ospf nsm 2583
debug ipv6 ospf packet 2584

© 2026 IP Infusion Inc. Proprietary

debug ipv6 ospf retransmission 2585
debug ipv6 ospf rib 2586
debug ipv6 ospf route 2587
default-information originate 2588
default-metric 2590
distance 2591
distribute bgp-ls 2592
distribute-list 2593
enable db-summary-opt 2595
ipv6 ospf authentication 2596
ipv6 ospf bfd 2597
ipv6 ospf cost 2598
ipv6 ospf dead-interval 2599
ipv6 ospf display route single-line 2600
ipv6 ospf hello-interval 2601
ipv6 ospf link-lsa-suppression 2602
ipv6 ospf mtu 2603
ipv6 ospf mtu-ignore 2604
ipv6 ospf neighbor 2605
ipv6 ospf network 2607
ipv6 ospf priority 2608
ipv6 ospf retransmit-interval 2609
ipv6 ospf transmit-delay 2610
ipv6 router ospf 2611
ipv6 te-metric 2613
log-adjacency-changes 2614
max-concurrent-dd 2615
passive-interface 2616
redistribute 2617
router-id 2619
router ipv6 ospf 2620
show debugging ipv6 ospf 2621
show ipv6 ospf 2622
show ipv6 ospf database 2624
show ipv6 ospf interface 2628
show ipv6 ospf neighbor 2630
show ipv6 ospf route 2633
show ipv6 ospf virtual-links 2635
show ipv6 ospfv3 topology 2637
show ipv6 route fast-reroute 2639

© 2026 IP Infusion Inc. Proprietary

show ipv6 vrf 2640
shutdown 2641
snmp restart ospf6 2642
summary-address 2643
timers spf exp 2645
OSPF VPN Commands 2646
capability vrf-lite 2647
domain-id 2648
router ospf vrf 2650
CSPF-TE Commands 2651
cspf default-retry-interval 2652
cspf disable-better-protection 2653
cspf enable-better-protection 2654
cspf enable-resource-constraint-crankback 2655
cspf tie-break 2656

© 2026 IP Infusion Inc. Proprietary

OSPFv2 Commands

This section describes the OSPFv2 commands.

area authentication 2386
area default-cost 2387
area filter-list 2388
area interface 2390
area interface authentication 2391
area interface cost 2392
area interface database-filter 2393
area interface dead-interval 2394
area interface hello-interval 2395
area interface IFNAME bfd 2396
area interface IFNAME priority 2397
area interface mpls ldp-igp 2398
area interface mpls ldp-igp sync ospf holddown-timer 2399
area interface network-type 2400
area interface passive 2401
area interface retransmit-interval 2402
area interface te-metric 2403
area nssa 2404
area range 2406
area sham-link 2408
area shortcut 2410
area stub 2411
area virtual-link 2412
auto-cost reference bandwidth 2414
bfd all-interfaces 2415
capability cspf 2416
capability lls 2417
capability opaque 2418
capability traffic-engineering 2419
capability vrf-lite 2420
clear ip ospf 2421
compatible rfc1583 2422
debug ip ospf dist-ls 2423
debug ip ospf lfa 2424
debug ip ospf redist 2425
debug ip ospf retransmission 2426

© 2026 IP Infusion Inc. Proprietary

debug ospf 2427
debug ospf database-timer rate-limit 2429
debug ospf events 2430
debug ospf ifsm 2431
debug ospf lsa 2432
debug ospf nfsm 2433
debug ospf nsm 2434
debug ospf packet 2435
debug ospf rib 2436
debug ospf route 2437
default-information originate 2438
default-metric 2440
distance 2441
distribute bgp-ls 2443
distribute-list 2444
enable db-summary-opt 2446
enable ext-ospf-multi-inst 2447
fast-reroute keep-all-paths 2448
fast-reroute per-prefix remote-lfa area 2449
fast-reroute per-prefix route-map 2450
fast-reroute terminate-hold-on interval 2451
fast-reroute tie-break 2452
host area 2454
ip ospf authentication 2455
ip ospf authentication-key 2456
ip ospf bfd 2458
ip ospf cost 2459
ip ospf database-filter 2460
ip ospf dead-interval 2461
ip ospf demand-circuit 2462
ip ospf disable 2463
ip ospf fast-reroute per-prefix candidate disable 2464
ip ospf flood-reduction 2465
ip ospf hello-interval 2466
ip ospf message-digest-key 2467
ip ospf mtu 2469
ip ospf mtu-ignore 2470
ip ospf multi-area 2471
ip ospf network 2473
ip ospf priority 2474

© 2026 IP Infusion Inc. Proprietary

ip ospf resync-timeout 2475
ip ospf retransmit-interval 2476
ip ospf transmit-delay 2477
log-adjacency-changes 2478
max-concurrent-dd 2479
maximum-area 2480
max-metric 2481
neighbor 2483
network 2484
ospf abr-type 2486
ospf area-interface-config-mode 2487
ospf flood-reduction 2488
ospf point-point rfc-incompatible 2489
ospf router-id 2490
overflow database 2491
overflow database external 2492
passive-interface 2493
redistribute 2495
router ospf 2497
show cspf lsp 2497
show cspf rsvp forwarding-timer 2500
show debugging ospf 2501
show ip ospf 2502
show ip ospf border-routers 2506
show ip ospf database brief 2508
show ip ospf database detail 2510
show ip ospf igp-shortcut-lsp 2518
show ip ospf igp-shortcut-route 2519
show ip ospf interface 2520
show ip ospf multi-area-adjacencies 2523
show ip ospf neighbor 2525
show ip ospf route 2529
show ip ospf sham-links 2532
show ip ospf valid 2534
show ip ospf virtual-links 2535
show ip protocols 2537
show ip route fast-reroute 2539
shutdown 2540
snmp context-name 2541
snmp restart ospf 2542

© 2026 IP Infusion Inc. Proprietary

summary-address 2543
te-metric 2545
timers lsa arrival 2546
timers spf exp 2547
timers throttle lsa 2548

© 2026 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 (page 2456) command to specify a simple text password.

Use the ip ospf message-digest-key (page 2467) 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
```

© 2026 IP Infusion Inc. Proprietary

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 (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
```

© 2026 IP Infusion Inc. Proprietary

```
(config)#ip access-list standard 1
(config)#deny host 172.22.0.0
(config)#router ospf 100
(config-router)#area 1 filter-list access 1 in
```

© 2026 IP Infusion Inc. Proprietary

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 version 4.2.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1
```

© 2026 IP Infusion Inc. Proprietary

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 version 4.2.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1 authentication null
```

© 2026 IP Infusion Inc. Proprietary

area interface 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 (108/ 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

```
area (<0-4294967295> |  A.B.C.D) interface <IFNAME> cost <1-65535>
no area (<0-4294967295> |  A.B.C.D) interface <IFNAME> cost
```

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

<0-65535>
Link-state metric

Default

By default, cost is 10.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1 cost 10
```

© 2026 IP Infusion Inc. Proprietary

area interface 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

```
area (A.B.C.D|<0-4294967295>) interface <IFNAME> database-filter all out
no area (A.B.C.D|<0-4294967295>) interface <IFNAME> database-filter all out
```

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

Disabled, all outgoing LSAs are flooded to the interface.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1 database-filter all out
```

© 2026 IP Infusion Inc. Proprietary

area interface 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

```
area (A.B.C.D|<0-4294967295>) interface <IFNAME> dead-interval <1-65535>
no area (A.B.C.D|<0-4294967295>) interface <IFNAME> dead-interval
```

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

dead-interval
Specify the interval

<1-65535>
Specify the interval in seconds

Default

No default value is specified.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1 dead-interval 10
```

© 2026 IP Infusion Inc. Proprietary

area interface 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

```
area (A.B.C.D|<0-4294967295>) interface <IFNAME> hello-interval <1-65535>
no area (A.B.C.D|<0-4294967295>) interface <IFNAME> hello-interval
```

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

hello-interval
Specify the interface name

<1-65535>
Specify the interval in seconds

Default

By default, hello interval is 10 seconds.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1 hello-interval 3
```

© 2026 IP Infusion Inc. Proprietary

area interface IFNAME 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

```
area (A.B.C.D|<0-4294967295>) interface <IFNAME> bfd (disable|)
no area (A.B.C.D|<0-4294967295>) interface <IFNAME> bfd (disable|)
```

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

disable
Disable the BFD option for neighbors on an interface

Default

No default value is specified.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1 bfd
```

© 2026 IP Infusion Inc. Proprietary

area interface IFNAME 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

```
area (<0-4294967295> |  A.B.C.D) interface <IFNAME> priority <0-255>
no area (<0-4294967295> |  A.B.C.D) interface <IFNAME> priority
```

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

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, priority is 1.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1 priority 3
```

© 2026 IP Infusion Inc. Proprietary

area interface mpls ldp-igp

Use this attribute to enable the MPLS LDP IGP SYNC on OSPF.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) interface <IFNAME> mpls ldp-igp sync ospf
no area (A.B.C.D|<0-4294967295>) interface <IFNAME> mpls ldp-igp sync ospf
```

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

sync
Enable MPLS LDP IGP SYNC on OSPF

Default

No default value is specified.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1 mpls ldp-igp sync ospf
```

© 2026 IP Infusion Inc. Proprietary

area interface mpls ldp-igp sync ospf holddown-timer

Use this attribute to configure the hold-down timer for which OSPF will wait for LDP to converge and advertises Max cost.

When the configured time expires, OSPF starts advertising the actual cost in the Router-LSA.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) interface <IFNAME> mpls ldp-igp sync ospf holddowntimer <1-2147483>
no area (A.B.C.D|<0-4294967295>) interface <IFNAME> mpls ldp-igp sync ospf
```

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

sync
Enable MPLS LDP IGP SYNC on OSPF

holddown-time
Holddown timer values in seconds.

<1-2147483>
OSPF area ID as a decimal value

Default

By default, holddown-timer is 70000 seconds.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1 mpls ldp-igp sync ospf holddown-timer 140000
```

© 2026 IP Infusion Inc. Proprietary

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 version 4.2.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1 network-type point-to-point
```

© 2026 IP Infusion Inc. Proprietary

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 version 4.2.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1 passive
```

© 2026 IP Infusion Inc. Proprietary

area interface 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

```
router ospf <0-65535> / area (A.B.C.D|<0-4294967295>) interface <IFNAME> retransmit-interval <1-3600>
router ospf <0-65535> / no area (A.B.C.D|<0-4294967295>) interface <IFNAME> retransmit-interval
```

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

retransmit-interval
Holddown timer values in seconds

<1-3600>
Specify the interval in seconds

Default

By default, retransmit interval is 5 seconds.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1 retransmit-interval 6
```

© 2026 IP Infusion Inc. Proprietary

area interface te-metric

This command sets the traffic engineering (TE) 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 ip ospf cost value for an interface is used in OSPF-TE Link State Advertisements.

Use the no parameter with this command to set the traffic engineering metric to its default.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) interface <IFNAME> te-metric <1-65535>
no area (A.B.C.D|<0-4294967295>) interface <IFNAME> te-metric <1-65535>
```

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

<1-65535>
Traffic engineering metric

Default

By default, the area interface cost value for an interface is used.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.3.0.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)#area 1 interface xe1 te-metric 6
```

© 2026 IP Infusion Inc. Proprietary

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|never)| stability-interval < 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.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 (advertise|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 .

advertise
Advertise this range.

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
```

© 2026 IP Infusion Inc. Proprietary

```
(config)#router ospf 100
(config-router)#no area 1 range 192.16.0.0/24
```

© 2026 IP Infusion Inc. Proprietary

area sham-link

Use this command to configure a sham link between two Provider Edge on MPLS backbone to make them appear as intra-area link.

If it is desired to have OSPF prefer the routes through the MPLS backbone over the routes through the backdoor link, then the routes through the MPLS backbone must be appear to be intra-area routes. To make a route through the backbone appear to be an intra-area route, it is necessary to make it appear as if there is an intra-area link connecting the two PE routers. This is what we refer to as a "sham link.

Use the no parameter with this command to remove a sham link.

Note: Sham-link is not supported on default VRF.

Command Syntax

```
area (A.B.C.D|<0-4294967295>) sham-link A.B.C.D A.B.C.D
area (A.B.C.D|<0-4294967295>) sham-link A.B.C.D A.B.C.D {cost <1-65535>}
no area (A.B.C.D|<0-4294967295>) sham-link A.B.C.D A.B.C.D
no area (A.B.C.D|<0-4294967295>) sham-link A.B.C.D A.B.C.D cost
```

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 source sham link.

A.B.C.D
Specify IP address of destination sham link.

<1-65535>
cost of sham link (cost preferably to be configured lower than the backdoor link cost).

Default

Cost: 1

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#configure terminal
```

© 2026 IP Infusion Inc. Proprietary

```
Enter configuration commands, one per line.  End with CNTL/Z.
(config)#router ospf 1 vrf1
(config-router)#area 1 sham-link 1.1.1.1 2.2.2.2 cost 10
(config-router)#
```

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.

Default

None

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 enable
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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|)|authentication-key LINE|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 (messagedigest|
null|)|authentication-key LINE|message-digest-key <1-255> md5 LINE}
no 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>}
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 message-digest
Enable cryptographic authentication on the configured virtual link.

authentication null
Enable null authentication on the configured virtual link.

authentication-key LINE
Set authentication key ID of 8 characters.

© 2026 IP Infusion Inc. Proprietary

message-digest-key <1-255>
Set message digest key.

md5 LINE
Specify the MD5 key.

dead-interval <1-65535>
The interval during which no packets are received and after which the router acknowledges a neighboring router as off-line. The interval in seconds. The default is 40 seconds.

hello-interval <1-65535>
The interval the router waits before it sends a hello packet. The interval in seconds. The default is 10 seconds.

retransmit-interval <1-3600>
The interval the router waits before it retransmits a packet. The interval in seconds. The default is 5 seconds.

transmit-delay <1-3600>
The interval the router waits before it transmits a packet. The interval in seconds. The default is 1 second

fall-over bfd
Specify fall-over Bidirectional Forwarding Detection (BFD) detection.

Default

Default intervals:

- Dead interval : 40 seconds
- Hello interval: 10 seconds
- Retransmit interval: 5 seconds
- Transmit delay: 1 second

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 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. 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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

capability cspf

Use this command to enable the Constrained Shortest Path First (CSPF) 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.

When CSPF is enabled, disable-better-protection option is by default enabled for OSPFv2.

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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.

Default

None

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ip ospf process
#clear ip ospf 555 process
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

debug ip ospf dist-ls

Use this command to display information about OSPF BGP link state distribution.

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

Command Syntax

```
debug ip ospf dist-ls
no debug ip ospf dist-ls
```

Parameters

None

Default

None

Command Mode

Privileged execution mode and Configure mode

Applicability

This command was introduced in OcNOS version 4.0.

Examples

```
# debug ip ospf dist-ls
```

© 2026 IP Infusion Inc. Proprietary

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

Default

None

Command Mode

Privileged execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ip ospf lfa
```

© 2026 IP Infusion Inc. Proprietary

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

Default

None

Command Mode

Privileged execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ip ospf redistribute detail
```

© 2026 IP Infusion Inc. Proprietary

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

Default

None

Command Mode

Privileged execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3

Examples

```
#debug ip ospf  retransmission
```

© 2026 IP Infusion Inc. Proprietary

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 (page 2429))

events
Debug OSPF events information (see debug ospf events (page 2430))

ifsm
Debug OSPF Interface State Machine (see debug ospf ifsm (page 2431))

lsa
Debug OSPF Link State Advertisement (see debug ospf lsa (page 2432))

nfsm
Debug OSPF Neighbor State Machine (see debug ospf nfsm (page 2433))

nsm
Debug OSPF NSM information (see debug ospf nsm (page 2434))

packet
Debug OSPF packets (see debug ospf packet (page 2435))

route
Debug OSPF route information (see debug ospf route (page 2437))

rib
Debug OSPF RIB information (see debug ospf rib (page 2436))

sr
Debug OSPF segment routing information

interface
Debug OSPF RIB interface

redistribute
Debug OSPF RIB redistribute

Default

None

© 2026 IP Infusion Inc. Proprietary

Command Mode

Privileged execution 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 2432) generate or debug ospf lsa (page 2432) 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

Default

None

Command Mode

Privileged execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
# debug ospf database-timer rate-limit
```

© 2026 IP Infusion Inc. Proprietary

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.

Default

None

Command Mode

Privileged execution 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
```

© 2026 IP Infusion Inc. Proprietary

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

Default

None

Command Mode

Privileged execution 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
```

© 2026 IP Infusion Inc. Proprietary

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.

Default

None

Command Mode

Privileged execution 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
```

© 2026 IP Infusion Inc. Proprietary

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

Default

None

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
```

© 2026 IP Infusion Inc. Proprietary

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.

Default

None

Command Mode

Privileged execution 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
```

© 2026 IP Infusion Inc. Proprietary

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.

Default

None

Command Mode

Privileged execution 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
```

© 2026 IP Infusion Inc. Proprietary

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.

Default

None

Command Mode

Privileged execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#debug ospf rib interface
#no debug ospf rib redistribute
```

© 2026 IP Infusion Inc. Proprietary

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.

Default

None

Command Mode

Privileged execution 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
```

© 2026 IP Infusion Inc. Proprietary

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 <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 1
The external link type associated with the default route advertised into the OSPF routing domain (see RFC 3101). Sets OSPF External Type 1 metric.

metric-type 1
The external link type associated with the default route advertised into the OSPF routing domain (see RFC 3101). Sets OSPF External Type 2 metric (default).

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.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 2495) 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
```

© 2026 IP Infusion Inc. Proprietary

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 <1-255>
Distance for all routes within an area

inter-area <1-255>
Distance for all routes from one area to another area.

external <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
```

© 2026 IP Infusion Inc. Proprietary

© 2026 IP Infusion Inc. Proprietary

distribute bgp-ls

Use this command to advertise IGP (OSPFv3) Updates to BGP module. Throttle value is optional, which can be used to delay the transmission of processing and sending updates to BGP

Use the `no` command to disable this function.

Command Syntax

```
distribute bgp-ls
distribute bgp-ls (throttle <val>)
no distribute bgp-ls
```

Parameters

`distribute`
distribute information

`bgp-ls`
Link State distribution to BGP

Default

By default, bgp-ls is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 6.1.0.

Examples

```
OcNOS#conf t
Enter configuration commands, one per line. End with CNTL/Z.
(config)#router ipv6 ospf 1
(config-router)#distribute bgp-ls throttle 10
```

© 2026 IP Infusion Inc. Proprietary

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 <1-65535>
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.

© 2026 IP Infusion Inc. Proprietary

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)#ip access-list list1 permit 172.10.0.0
(config)#router ospf 100
(config-router)#distribute-list list1 out bgp
(config-router)#redistribute bgp
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

enable ext-ospf-multi-inst

Use this command to run multiple OSPF instances on a subnet.

Use the no parameter with this command to disable OSPF multiple-instance support and reset all OSPF instances to the default instance ID.

Command Syntax

```
enable ext-ospf-multi-inst
no enable ext-ospf-multi-inst
```

Parameters

None

Default

Multiple-instance support is disabled. The default instance ID is 0.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)# enable ext-ospf-multi-inst
```

© 2026 IP Infusion Inc. Proprietary

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

Default

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
```

© 2026 IP Infusion Inc. Proprietary

fast-reroute per-prefix remote-lfa area

Use this command to configure remote-lfa on an area in OSPF.

Use the no form of this command remove the remote-lfa configuration.

Command Syntax

```
fast-reroute per-prefix remote-lfa area (A.B.C.D|<0-4294967295>) tunnel mpls-ldp
no fast-reroute per-prefix remote-lfa area (A.B.C.D|<0-4294967295>) tunnel mpls-ldp
```

Parameters

A.B.C.D
Area-ID in IP address notation.

0-4294967295
Area-ID in integer notation.

Default

None

Command Mode

Router OSPF mode.

Applicability

This command was introduced before OcNOS version 6.0.0.

Examples

```
(config)#router ospf 1
(config-router)#fast-reroute per-prefix remote-lfa area 0 tunnel mpls-ldp
(config-router)#

(config-router)#no fast-reroute per-prefix remote-lfa area 0 tunnel mpls-ldp
(config-router)#commit
```

© 2026 IP Infusion Inc. Proprietary

fast-reroute per-prefix route-map

Use this command to configure route-map in OSPF for fast-reroute calculation.

Use the no form of this command remove the route-map applied on fast-reroute calculation.

Command Syntax

```
fast-reroute per-prefix route-map WORD
no fast-reroute per-prefix route-map
```

Parameters

WORD
Route-map name.

Default

None

Command Mode

Router OSPF mode.

Applicability

This command was introduced before OcNOS version 6.0.0.

Examples

```
(config)#router ospf 1
(config-router)#fast-reroute per-prefix route-map rmap1
(config-router)#

(config-router)#no fast-reroute per-prefix route-map
(config-router)#commit
```

© 2026 IP Infusion Inc. Proprietary

fast-reroute terminate-hold-on interval

Use this command to set the delay of primary route installation (to avoid micro loop) after a failover.

Use the no form of this command to set the termination hold-on timer to its default value (600 milliseconds).

Command Syntax

```
fast-reroute terminate-hold-on interval <100-100000>
no fast-reroute terminate-hold-on interval
```

Parameters

<100-100000>
Hold on interval in milliseconds

Default

600 milliseconds

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 1
(config-router)#fast-reroute terminate-hold-on interval 7000
(config-router)#no fast-reroute terminate-hold-on interval
```

© 2026 IP Infusion Inc. Proprietary

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 as shown in Table 69 table.

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|downstream-path|secondary-path) index <1-255>
no fast-reroute tie-break
no fast-reroute tie-break (primary-path|interface-disjoint|node-protecting|broadcast-interface-
disjoint|downstream-path|secondary-path)
```

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
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.

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.

secondary-path
Prefer a non-ECMP backup path.

index <1-255>
Tie break priority. A lower value has higher preference. Range of priority values.

© 2026 IP Infusion Inc. Proprietary

Default

Table 69. Default Value

`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 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 A.B.C.D
Set the OSPF area ID in IPv4 address format.

area <0-4294967295>
Set the OSPF Area ID as a decimal value.

cost <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
```

© 2026 IP Infusion Inc. Proprietary

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 xe2 is configured to have no authentication. This will override any text or MD5 authentication configured on this interface.

```
#configure terminal
(config)#interface xe2
(config-if)#ip ospf authentication null
```

© 2026 IP Infusion Inc. Proprietary

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 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 xe2 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 xe2
(config-if)#ip ospf 3.3.3.3 authentication-key test

#sh run int xe2
!
interface xe2
```

© 2026 IP Infusion Inc. Proprietary

```
ip ospf authentication
ip ospf authentication-key 0x94eebee8c349a4b0
!
```

Encrypted password of 16 characters which can be obtained from `sh run <int>` command.

```
(config)#int xe2
(config-if)#ip ospf authentication-key 0x94eebee8c349a4b0
```

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 xe1
(config-if)#ip ospf bfd
```

© 2026 IP Infusion Inc. Proprietary

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>
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 xe1 for IP address 10.10.10.50.

```
#configure terminal
(config)#interface xe1
(config-if)#ip ospf 10.10.10.50 cost 10
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 <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 xe1 interface.

```
#configure terminal
(config)#interface xe1
(config-if)#ip ospf dead-interval 10
```

© 2026 IP Infusion Inc. Proprietary

ip ospf demand-circuit

Use this command to enable hello suppression and LSA suppression sent on OSPF interface. Enabling demand circuit on one interface enables hello suppression only for that particular interface.

Note: Hellos and LSAs are suppressed on Point-to-point and Point-to-multipoint links and only LSAs are suppressed for a broadcast link.

Use the no parameter with this command to disable hello suppression and LSA suppression

Command Syntax

```
ip ospf demand-circuit
no ip ospf demand-circuit
```

Parameters

None

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows configuring demand-circuit on xe1 interface.

```
#configure terminal
(config)#interface xe1
(config-if)#ip ospf demand-circuit

(config-if)#no ip ospf demand-circuit
```

© 2026 IP Infusion Inc. Proprietary

ip ospf disable

Use this command to completely disable OSPF packet processing on an interface.

This command overrides the network (page 2484) 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 xe1
(config-if)#ip ospf disable all
```

© 2026 IP Infusion Inc. Proprietary

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 xe1
(config-if)# ip ospf fast-reroute per-prefix candidate disable
```

© 2026 IP Infusion Inc. Proprietary

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 xe1
(config-if)#ip ospf flood-reduction

(config-if)#no ip ospf flood-reduction
```

© 2026 IP Infusion Inc. Proprietary

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 <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 xe1.

```
#configure terminal
(config)#interface xe1
(config-if)#ip ospf hello-interval 3
```

© 2026 IP Infusion Inc. Proprietary

ip ospf message-digest-key

Use this command to register an MD5 key for OSPF authentication.

- 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.

Use the no parameter with this command to remove an MD5 key.

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 <1-255>
Specify a message digest authentication password (key ID).

md5 WORD
Specify an encrypted password (key) (1-16 characters).

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. Encrypted password of 16 characters which can be obtained from `sh run <int>` command.

© 2026 IP Infusion Inc. Proprietary

```
#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
```

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
```

© 2026 IP Infusion Inc. Proprietary

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 <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 xe1
(config-if)#ip ospf mtu 1480
```

© 2026 IP Infusion Inc. Proprietary

ip ospf mtu-ignore

Use this command to configure OSPF so that it does not check the MTU size during Database Description (DD) 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 xe1
(config-router)#ip ospf mtu-ignore
```

© 2026 IP Infusion Inc. Proprietary

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

<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.

neighbor 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 xe1
(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
```

© 2026 IP Infusion Inc. Proprietary

```
(config)#interface xe1
(config-if)#ip ospf 0 multi-area 1
(config)#interface xe2
(config-if)#ip ospf 0 multi-area 1 neighbor 5.5.5.2
(config-if)# no ip ospf 0 multi-area 1
```

© 2026 IP Infusion Inc. Proprietary

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.

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 xe1 interface.

```
#configure terminal
(config)#interface xe1
(config-if)#ip ospf network point-to-point
```

© 2026 IP Infusion Inc. Proprietary

ip ospf priority

Use this command to set the router priority to determine the designated router (DR) for the network.

Notes:

- 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 <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 xe1 interface.

```
#configure terminal
(config)#interface xe1
(config-if)#ip ospf priority 3
```

© 2026 IP Infusion Inc. Proprietary

ip ospf resync-timeout

Use this command to set the interval after which adjacency is reset if out-of-band re-synchronization has not occurred. The interval period starts from the time a restart signal is received from a neighbor.

Use the no parameter with this command to set the default value (40 seconds).

Command Syntax

```
ip ospf (A.B.C.D|) resync-timeout <1-65535>
no ip ospf (A.B.C.D|) resync-timeout
```

Parameters

A.B.C.D
IP address of the interface.

<1-65535>
Re-synchronization timeout of the interface in seconds.

Default

By default, the re-synchronization timeout is 40 seconds.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows setting the OSPF re-synchronization timeout value to 65 seconds on the xe1 interface.

```
#configure terminal
(config)#interface xe1
(config-if)#ip ospf resync-timeout 65
```

© 2026 IP Infusion Inc. Proprietary

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 <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 xe1 interface.

```
#configure terminal
(config)#interface xe1
(config-if)#ip ospf retransmit-interval 6
```

© 2026 IP Infusion Inc. Proprietary

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 (1 second).

Command Syntax

```
ip ospf (A.B.C.D|) transmit-delay <1-3600>
no ip ospf (A.B.C.D|) transmit-delay
```

Parameters

A.B.C.D
IPv4 address of the interface.

<1-3600>
Time in seconds to transmit a link-state update.

Default

By default, the 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 xe1 interface.

```
#configure terminal
(config)#interface xe1
(config-if)#ip ospf transmit-delay 3
```

© 2026 IP Infusion Inc. Proprietary

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

If not specified, brief is the default option.

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 <1-16777215 >
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 <5-86400>
Startup metric values for router LSA

summary-lsa <1-16777215>
Summary LSA max metric values

Default

None

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
```

© 2026 IP Infusion Inc. Proprietary

```
!
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
!
```

© 2026 IP Infusion Inc. Proprietary

neighbor

Use this command to configure OSPF routers interconnecting to Non-Broadcast Multi-Access (NBMA) 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
Interface IP address of the neighbor.

priority <0-255>
Router priority of the non-broadcast neighbor associated with the specified IP address.This parameter does not apply to point-to-multipoint interfaces.

poll-interval <1-2147483647>
Reduced rate at which routers continue to send hello packets when a neighboring router has become inactive. Dead neighbor polling interval in seconds. Set this value much larger than hello interval.

cost <1-65535>
OSPF cost metric for point-to-multipoint neighbor.

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
```

© 2026 IP Infusion Inc. Proprietary

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 A.B.C.D
OSPF area ID in IPv4 address format.

area <0-4294967295>
OSPF area ID as a decimal value.

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.

© 2026 IP Infusion Inc. Proprietary

```
#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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 version 4.2.

Examples

```
#configure terminal
(config)#ospf area-interface-config-mode
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

ospf point-point rfc-incompatible

Use this command to disable RFC 2328 compatibility on point-to-point interfaces.

By default, in a point-to-point network OSPF packets are transmitted using multicast addresses. By disabling point-to-point RFC 2328 compatibility, packets are transmitted using the unicast address of the neighbor.

Use the no parameter with this command to enable RFC 2328 compatibility on point-to-point interfaces.

Command Syntax

```
ospf point-point rfc-incompatible
no ospf point-point rfc-incompatible
```

Parameters

None

Default

By default, RFC 2328 compatibility is enabled (packets transmitted using multicast addresses).

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ospf 100
(config-router)# ospf point-point rfc-incompatible
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 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 configures all interfaces as passive. If it is re-applied, the other passive configurations will be removed and all interfaces will become passive again.

Command Syntax

```
passive-interface
passive-interface IFNAME ( ( disable | 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

disable
Disable 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
(config-router)#passive-interface xe1 enable
```

© 2026 IP Infusion Inc. Proprietary

```
(config-router)#passive-interface xe2 10.1.2.3 enable
```

redistribute

This command redistributes routes from routing protocols, static routes, and kernel routes into an OSPF routing table.

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

Command Syntax

```
redistribute (connected|static|rip|bgp|isis (WORD|)|ospf (<1-65535>|)) {metric <0-16777214>|metric-
type (1|2)|route-map WORD|tag <0-4294967295>}
no redistribute (connected|static|rip|bgp|isis(WORD|)|ospf (<1-65535>|)) metric|metric-type|route-
map|tag}
```

Parameters

connected
Redistribute connected routes.

static
Redistribute static routes.

rip
Redistribute RIP routes.

bgp
Redistribute BGP routes.

isis
Redistribute IS-IS routes.

isis WORD
IS-IS instance identifier.

ospf
Redistribute OSPF routes.

ospf <1-65535>
OSPF process ID to redistribute a particular OSPF instance into another OSPF instance. If not specified, this command redistributes into the OSPF instance with process ID 0.

metric <0-16777214>
External metric value.

metric-type 1
OSPF External Type 1 metrics. (see RFC 3101)

metric-type 2
OSPF External Type 2 metrics. (see RFC 3101)

route-map WORD
Route map name.

tag <0-4294967295>
Route tag value to use as a “match” value for controlling redistribution via route maps.

Default

By default, redistribution is disabled.

© 2026 IP Infusion Inc. Proprietary

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.

```
(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.

```
(config)#router ospf 1
(config-router)#redistribute ospf 2 metric 10 metric-type 1 route-map rmp1 tag 3
```

© 2026 IP Infusion Inc. Proprietary

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 RSVP 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 <0-65535>

no router ospf
no router ospf <0-65535>
```

Parameters

<0-65535>
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 cspf lsp

Use this command to display CSPF statistics, including counts or status of CSPF-computed paths for label-switched paths. Displays the actual CSPF-selected paths with extensive constraint detail.

© 2026 IP Infusion Inc. Proprietary

Command Syntax

```
Show cspf lsp
Show cspf lsp SESSION-NAME
show cspf lsp count
```

Parameters

SESSION-NAME

Use this parameter to display only the CSPF LSPs matching a specific RSVP session name.

count

Use this parameter to display the LSP count.

Command Mode

Privileged execution mode and Execution mode

Applicability

This command was introduced before OcNOS version 7.0.0.

Example

Following is a sample output from the command.

```
OcNOS#show cspf lsp
   LSP Name        : t1-Primary
   Trunk ID        : 5001
   Lsp Id          : 2202
   OSPF ID         : 1
   Ingress         : 2.2.2.2
   Egress          : 3.3.3.3
   Ext Tunnel ID   : 2.2.2.2
   LSP Type        : 0
   Client ID       : 1
   State           : Established(3)
   Setup Priority  : 7
   Hold Priority   : 0
   Hop Limit       : 255
   Include Mask    : 0x0
   Exclude Mask    : 0x0
   Include All Mask: 0x0
   LSP Metric      : 1
   LSP Reoptimize  : Yes
   LSP in ReopList : Yes
   LSP ReopTime    : 1 minutes (Offset 18 seconds)
   Next Reop Due in: 60 seconds
   Computed ERO    :
       30.0.0.3

   LSP Name        : BL-30.0.0.3-3.3.3.3-100-Bypass
   Trunk ID        : 5002
   Lsp Id          : 2201
   OSPF ID         : 1
   Ingress         : 2.2.2.2
   Egress          : 3.3.3.3
   Ext Tunnel ID   : 2.2.2.2
   LSP Type        : 2
   Client ID       : 1
   State           : Established(3)
```

© 2026 IP Infusion Inc. Proprietary

```
   Setup Priority  : 7
   Hold Priority   : 0
   Hop Limit       : 255
   Include Mask    : 0x0
   Exclude Mask    : 0x0
   Include All Mask: 0x0
   LSP Metric      : 1
   LSP Reoptimize  : Yes
   LSP in ReopList : Yes
   LSP ReopTime    : 1 minutes (Offset 13 seconds)
   Next Reop Due in: 55 seconds
   Exclude Path Constraint :
       30.0.0.3    exclude link
   Computed ERO    :
       37.0.0.3

OcNOS#
```

Following is a sample output from the command using the `SESSION-NAME` parameter.

```
OcNOS#show cspf lsp BL
 LSP detail
   LSP Name               : BL-30.0.0.3-3.3.3.3-100-Bypass
   Trunk ID               : 5002
   LSP ID                 : 2202
   ISIS Tag Name          : 10
   Client ID              : 5
   State                  : Established(3)
   Ingress                : 2.2.2.2
   Egress                 : 3.3.3.3
   Ext Tunnel ID          : 2.2.2.2
   LSP Type               : 2
   Setup Priority         : 7
   Hold Priority          : 0
   Hop Limit              : 255
   Include Mask           : 0x0
   Exclude Mask           : 0x0
   Include All Mask       : 0x0
   LSP Metric             : 10
   LSP Reoptimize         : Yes
   LSP in Reop-List       : Yes
   LSP ReopTime           : 1 minutes (Offset 14 seconds)
   Next Reop Due in       : 23 seconds
   Computed ERO           :
       37.0.0.3

OcNOS#
```

Following is a sample output from the command using the `count` parameter.

```
OcNOS#show cspf lsp count
 ISIS cspf lsp count
  LSP Count for Setup priority 0     : 0
  LSP Count for Setup priority 1     : 0
  LSP Count for Setup priority 2     : 0
  LSP Count for Setup priority 3     : 0
  LSP Count for Setup priority 4     : 0
  LSP Count for Setup priority 5     : 0
  LSP Count for Setup priority 6     : 0
  LSP Count for Setup priority 7     : 2
OcNOS#
```

© 2026 IP Infusion Inc. Proprietary

show cspf rsvp forwarding-timer

This command displays the information of Graceful Restart capable RSVP client to ISIS or OSPF, CSPF that are currently shutdown.

Command Syntax

```
show cspf rsvp forwarding-timer
```

Parameters

None

Default

None

Command Mode

Privileged Exec modes

Applicability

This command was introduced before OcNOS version 5.0.

Example

```
#show cspf rsvp forwarding-timer
CSPF Server Protocol-Name GR-State Time Remaining (sec) Disconnected-time
ISIS RSVP ACTIVE 88 2021/08/18 04:49:23
```

© 2026 IP Infusion Inc. Proprietary

show debugging ospf

Use this command to display the set OSPF debugging option.

Command Syntax

```
show debugging ospf
```

Parameters

None

Default

None

Command Mode

Privileged execution 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
```

© 2026 IP Infusion Inc. Proprietary

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.

Default

None

Command Mode

Privileged execution 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
```

Here is the explanation of the "show command" output fields.

Table 70. 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.
Connected to MPLS VPN Super backbone	VRF is enabled and the process is connected to the MPLS VPN backbone.
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.

© 2026 IP Infusion Inc. Proprietary

Table 70. 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.
Area has message digest authentication	Area uses MD5 authentication.
SPF algorithm last executed	As stated.

© 2026 IP Infusion Inc. Proprietary

Table 70. show ip ospf output details (continued)

Field	Description
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.

Default

None

Command Mode

Privileged execution 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
```

Here is the explanation of the "show command" output fields.

Table 71. 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.

© 2026 IP Infusion Inc. Proprietary

Table 71. border router output details (continued)

Field	Description
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

Default

None

Command Mode

Privileged execution 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
```

Here is the explanation of the "show command" output fields.

© 2026 IP Infusion Inc. Proprietary

Table 72. 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.

© 2026 IP Infusion Inc. Proprietary

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.

Default

None

© 2026 IP Infusion Inc. Proprietary

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

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
```

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])
```

© 2026 IP Infusion Inc. Proprietary

```
  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)
```

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 for Flood Reduction

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
```

Output Fields Description

Database Detail Header Fields

Here is the explanation of the "show command" output fields.

Table 73. 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 74.
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:

Table 73. ospf database detail header fields (continued)

Field	Description
	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
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

Here is the explanation of the "show command" output fields.

Table 74. ospf LSA option bits output details

Bit	Description
DN	Used in MPLS-based L3 VPNs. When a route learned from a customer network via OSPF is advertised across a BGP/MPLS 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.

© 2026 IP Infusion Inc. Proprietary

Table 74. ospf LSA option bits output details (continued)

Bit	Description
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)
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)

Here is the explanation of the "show command" output fields.

Table 75. 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

Table 75. router LSAs (continued)

Field	Description
(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)

Here is the explanation of the "show command" output fields.

Table 76. 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)

Here is the explanation of the "show command" output fields.

Table 77. 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)

Here is the explanation of the "show command" output fields.

Table 78. 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	Data traffic for the advertised destination is forwarded to this address.

© 2026 IP Infusion Inc. Proprietary

Table 78. external LSAs (continued)

Field	Description
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)

Here is the explanation of the "show command" output fields.

Table 79. 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.

© 2026 IP Infusion Inc. Proprietary

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

Default

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
```

Here is the explanation of the "show command" output fields.

Table 80. 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.

© 2026 IP Infusion Inc. Proprietary

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.

Default

None

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
```

Here is the explanation of the "show command" output fields.

Table 81. 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.

© 2026 IP Infusion Inc. Proprietary

show ip ospf interface

Use this command to display interface information for OSPF.

Command Syntax

```
show ip ospf interface (IFNAME|) brief
show ip ospf interface brief
```

Parameters

IFNAME
Interface name.

brief
Brief summary of interface

Default

None

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
    LDP-OSPF Sync configured
    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
```

Here is the explanation of the "show command" output fields.

© 2026 IP Infusion Inc. Proprietary

Table 82. 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.
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 interface brief

Interface    PID   Area            Intf ID    Cost  State           Neighbors  Status
lo           1     0.0.0.0         1          1     Loopback        0            Up
Interface    PID   Area            Intf ID    Cost  State           Neighbors  Status
xe1          1     0.0.0.0         4          10    DR              1            Up
Interface    PID   Area            Intf ID    Cost  State           Neighbors  Status
xe2          1     0.0.0.0         14         15    DR              1            Up
Interface    PID   Area            Intf ID    Cost  State           Neighbors  Status
xe3          1     0.0.0.0         15         20    DR              1            Up
```

© 2026 IP Infusion Inc. Proprietary

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.

Default

None

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
```

Here is the explanation of the "show command" output fields.

Table 83. 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

© 2026 IP Infusion Inc. Proprietary

Table 83. show ip ospf multi-area-adjacencies output details (continued)

Field	Description
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.
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	

© 2026 IP Infusion Inc. Proprietary

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

Default

None

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

Here is the explanation of the "show command" output fields.

© 2026 IP Infusion Inc. Proprietary

Table 84. OSPF neighbor output details

Field	Description
OSPF process	OSPF process identifier.
Neighbor ID	OSPF router identifier of the neighbor.
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

Here is the explanation of the "show command" output fields.

Table 85. 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 84.
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 74.
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	Off if hello suppression is enabled, on otherwise.

© 2026 IP Infusion Inc. Proprietary

Table 85. OSPF neighbor output detail (continued)

Field	Description
Description Retransmission	
Thread Link State Request Retransmission	Off if hello suppression is enabled, on otherwise.
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.

Default

None

Command Mode

Privileged execution 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 Fields

Each entry in this table has a code preceding it indicating the source of the routing entry.

Here is the explanation of the "show command" output fields.

Table 86. 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

Here is the explanation of the "show command" output fields.

Table 87. route entry output details

Field	Description
Codes	As explained in Table 86.
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.

© 2026 IP Infusion Inc. Proprietary

```
#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
```

© 2026 IP Infusion Inc. Proprietary

show ip ospf sham-links

Use this command to display sham link information.

Command Syntax

```
show ip ospf (<0-65535>|) sham-links (brief|)
```

Parameters

<0-65535>
The ID of the router process for which information will be displayed.

brief
Display summary of OSPF sham-links.

Default

None

Command Mode

Privileged Exec mode and Exec mode

Applicability

This command was introduced before OcNOS version 5.0.

Examples

The following is the display of the sham link information for two routers, one with the sham link up and one with sham link down.

```
#show ip ospf sham-links
Sham Link SLINK0 to destination 2.2.2.2 is down
  area 0.0.0.1 source 1.1.1.1
  Transmit Delay is 1 sec, State Down,
  Timer intervals , Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in inactive
  No authentication
    Adjacency state Down

#show ip ospf sham-links brief

  OSPF Process ID 1 VRF vrf1
  Total number of slinks: 1
  Shamlink Name    DestinationID  SourceID  Area    Status
  SLINK0          2.2.2.2              1.1.1.1  0.0.0.1  down

  OSPF Process ID 100 VRF default
  Total number of slinks: 0
  Shamlink Name    DestinationID  SourceID  Area    Status

#show ip ospf sham-links
Sham Link SLINK0 to destination 44.44.44.44 is up
  area 0.0.0.2 source 33.33.33.33
  Transmit Delay is 1 sec, State Point-To-Point,
  Timer intervals , Hello 10, Dead 40, Wait 40, Retransmit 5
```

© 2026 IP Infusion Inc. Proprietary

```
    Hello due in 00:00:05
  No authentication
    Adjacency state Full

#show ip ospf sham-links brief

  OSPF Process ID 1 VRF vrf1
  Total number of slinks: 1
  Shamlink Name    DestinationID  SourceID  Area    Status
  SLINK0          44.44.44.44          33.33.33.33 0.0.0.2  up

  OSPF Process ID 100 VRF default
  Total number of slinks: 0
  Shamlink Name    DestinationID  SourceID  Area    Status
```

Here is the explanation of the "show command" output fields.

Table 88. show ip ospf sham-links output details

Field	Description
sham Link	Sham link name, the Destination to which it is connected, and the state of the link.
Area	The area the sham link is connected to
Source	Source address of the sham link
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	The Default 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.

© 2026 IP Infusion Inc. Proprietary

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.

Default

None

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show ip ospf 1 opaque-link valid
```

© 2026 IP Infusion Inc. Proprietary

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.

Default

None

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.

```
OcNOS#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

```
OcNOS#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
```

© 2026 IP Infusion Inc. Proprietary

```
 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

```
OcNOS#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
```

Here is the explanation of the "show command" output fields.

Table 89. 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.

© 2026 IP Infusion Inc. Proprietary

show ip protocols

Use this command to display OSPF process parameters and statistics.

Command Syntax

```
show ip protocols
show ip protocols ospf
```

Parameters

None

Default

None

Command Mode

Privileged execution 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
```

Here is the explanation of the "show command" output fields.

Table 90. 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.

© 2026 IP Infusion Inc. Proprietary

Table 90. show ip protocols output details (continued)

Field	Description
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.
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.

© 2026 IP Infusion Inc. Proprietary

show ip route fast-reroute

Use this command to display routes with alternate next hops.

Command Syntax

```
show ip route fast-reroute
```

Parameters

None

Default

None

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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.

Default

None

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 to restart SNMP in OSPF.

Command Syntax

```
snmp restart ospf
```

Parameter

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 ospf
```

© 2026 IP Infusion Inc. Proprietary

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 <0-4294967295>
Set tag value to use as a “match” value for controlling redistribution via route maps.

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
```

© 2026 IP Infusion Inc. Proprietary

```
(config-router)#summary-address 172.16.0.0/16 tag 3
```

te-metric

This command sets the traffic engineering (TE) 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 ip ospf cost (page 2459) value for an interface is used in OSPF-TE Link State Advertisements.

Use the no parameter with this command to set the traffic engineering metric to its default.

Command Syntax

```
te-metric <1-65535>
no te-metric
```

Parameters

<1-65535>
Traffic engineering metric.

Default

By default, the ip ospf cost (page 2459) value for an interface is used.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface xe1
(config-if)#te-metric 6
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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.

<1-600000>

Hold interval: The hold time in milliseconds. This value is used to calculate the subsequent rate limiting times for LSA generation.

<1-600000>

Maximum interval: The maximum wait time in milliseconds between generation of the same LSA.

Default

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
```

© 2026 IP Infusion Inc. Proprietary

OSPFv2 Graceful Restart Commands

This section describes the OSPFv2 graceful restart commands.

action interface 2550
capability restart 2551
event nbr 2552
debug ip ospf graceful-restart 2553
if-event-track ospf event 2554
ospf restart grace-period 2555
ospf restart helper 2556
restart ip ospf graceful 2557

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 <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 ip 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
```

© 2026 IP Infusion Inc. Proprietary

OSPFv3 Graceful Restart Commands

This section describes the OSPFv3 graceful restart commands.

capability restart 2559
ipv6 ospf restart grace-period 2560
ipv6 ospf restart helper 2561
ipv6 ospf restart planned-only 2562
restart ipv6 ospf graceful 2563

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

OSPFv3 Commands

This section provides an alphabetized reference for each of the OSPFv3 commands. It includes the following commands:

abr-type 2566
area default-cost 2567
area nssa 2568
area range 2570
area stub 2571
area virtual-link 2572
auto-cost reference bandwidth 2574
bfd all-interfaces 2575
clear ipv6 ospf process 2576
debug ipv6 ospf 2577
debug ipv6 ospf bfd 2578
debug ipv6 ospf events 2579
debug ipv6 ospf ifsm 2580
debug ipv6 ospf lsa 2581
debug ipv6 ospf nfsm 2582
debug ipv6 ospf nsm 2583
debug ipv6 ospf packet 2584
debug ipv6 ospf retransmission 2585
debug ipv6 ospf rib 2586
debug ipv6 ospf route 2587
default-information originate 2588
default-metric 2590
distance 2591
distribute bgp-ls 2592
distribute-list 2593
enable db-summary-opt 2595
ipv6 ospf authentication 2596
ipv6 ospf bfd 2597
ipv6 ospf cost 2598
ipv6 ospf dead-interval 2599
ipv6 ospf display route single-line 2600
ipv6 ospf hello-interval 2601
ipv6 ospf link-lsa-suppression 2602
ipv6 ospf mtu 2603
ipv6 ospf mtu-ignore 2604

© 2026 IP Infusion Inc. Proprietary

ipv6 ospf neighbor 2605
ipv6 ospf network 2607
ipv6 ospf priority 2608
ipv6 ospf retransmit-interval 2609
ipv6 ospf transmit-delay 2610
ipv6 router ospf 2611
ipv6 te-metric 2613
log-adjacency-changes 2614
max-concurrent-dd 2615
passive-interface 2616
redistribute 2617
router-id 2619
router ipv6 ospf 2620
show debugging ipv6 ospf 2621
show ipv6 ospf 2622
show ipv6 ospf database 2624
show ipv6 ospf interface 2628
show ipv6 ospf neighbor 2630
show ipv6 ospf route 2633
show ipv6 ospf virtual-links 2635
show ipv6 ospfv3 topology 2637
show ipv6 route fast-reroute 2639
show ipv6 vrf 2640
shutdown 2641
snmp restart ospf6 2642
summary-address 2643
timers spf exp 2645

© 2026 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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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.

<0-2147483647>
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.

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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

None

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
```

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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

Hello interval:10 seconds.

Dead interval:40 seconds.

Retransmit interval: 5 seconds.

Transmit delay: 1 second

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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 execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear ipv6 ospf Tag1 process
```

© 2026 IP Infusion Inc. Proprietary

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 (page 2578))

events
Debug OSPFv3 events (see debug ipv6 ospf events (page 2579)).

ifsm
Debug OSPFv3 Interface State Machines (see debug ipv6 ospf ifsm (page 2580)).

lsa
Debug OSPFv3 Link State Advertisements (see debug ipv6 ospf lsa (page 2581)).

nfsm
Debug OSPFv3 Neighbor State Machines (see debug ipv6 ospf nfsm (page 2582)).

nsm
Debug OSPFv3 NSM information (see debug ipv6 ospf nsm (page 2583)).

packet
Debug OSPFv3 packets (see debug ipv6 ospf packet (page 2584)).

retransmission
Debug OSPFv3 retransmission information. (see debug ipv6 ospf retransmission (page 2585))

rib
Debug OSPFv3 Routing Information Base.(see debug ipv6 ospf rib (page 2586))

route
Debug OSPFv3 route information (see debug ipv6 ospf route (page 2587)).

Command Mode

Privileged execution modeand Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf all
```

© 2026 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 execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf bfd
```

© 2026 IP Infusion Inc. Proprietary

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 execution 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
```

© 2026 IP Infusion Inc. Proprietary

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 execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf ifsm status
```

© 2026 IP Infusion Inc. Proprietary

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 execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf lsa
```

© 2026 IP Infusion Inc. Proprietary

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 execution modeand Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf nfsm events
#no debug ipv6 ospf nfsm timers
```

© 2026 IP Infusion Inc. Proprietary

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 execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf nsm interface
```

© 2026 IP Infusion Inc. Proprietary

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 execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf packet ls-request
```

© 2026 IP Infusion Inc. Proprietary

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 execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf  retransmission
```

© 2026 IP Infusion Inc. Proprietary

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 execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug ipv6 ospf rib interface
```

© 2026 IP Infusion Inc. Proprietary

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 execution modeand Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#no debug ipv6 ospf route
#debug ipv6 ospf route ia
```

© 2026 IP Infusion Inc. Proprietary

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).

2 sets OSPF External Type 2 metric

Command Mode

Router mode

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 2617) 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

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 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> <1-255> <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
```

© 2026 IP Infusion Inc. Proprietary

distribute bgp-ls

Use this command to advertise IGP (OSPFv3) Updates to BGP module. Throttle value is optional, which can be used to delay the transmission of processing and sending updates to BGP

Use the no command to disable this function.

Command Syntax

```
distribute bgp-ls
distribute bgp-ls (throttle <val>)
no distribute bgp-ls
```

Parameters

distribute
Distribute information.

bgp-ls
Link State distribution to BGP.

Default

Disable

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 6.6.1.

Examples

```
OcNOS#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
(config)#router ipv6 ospf 1
(config-router)#distribute bgp-ls throttle 10
```

© 2026 IP Infusion Inc. Proprietary

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

None

Command Mode

Router mode

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

None

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS version 6.0.0.

Example

```
#configure terminal
(config)#interface eth1
(config-if)#ipv6 ospf authentication cryptomap map1
```

© 2026 IP Infusion Inc. Proprietary

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

Disabled

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

ipv6 ospf display route single-line

Use this command to display the output of the show ipv6 ospf route (page 2633) 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

show ipv6 ospf route (page 2633) 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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

ipv6 ospf mtu

Use this command to set MTU size for OSPFv3 to construct packets based on this value. Whenever OSPFv3 constructs packets, it uses interface MTU size as Maximum IP packet size. This command forces OSPFv3 to use the specified value overriding the actual interface MTU size.

Use the `no` parameter with this command to return to the default value.

Command syntax

```
ipv6 ospf mtu <1280-65535> instance-id (<0-31>|<64-95>)
no ipv6 ospf mtu instance-id (<0-31>|<64-95>)
```

Parameters

<1280-65535>
Interface MTU value from OSPFv3's perspective.

<0-31>
Interface instance ID for IPv6 unicast

<64-95>
Interface instance ID for IPv4 unicast.

Default

OSPFv3 uses interface MTU derived from the kernel and Instance ID as 0 if not configured.

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS version 4.0.

Examples

```
#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
(config)#interface eth1
(config-if)#ipv6 ospf mtu 1287 instance-id 10
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

Cost is 10.

Poll interval is 120 seconds.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

1

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#ipv6 ospf priority 127
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

Disabled

Command Mode

Interface mode

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 2598) 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

Zero

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 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 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

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
```

© 2026 IP Infusion Inc. Proprietary

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 (kernel|connected|static|rip|bgp|isis|ospf (WORD|<1-65535>|)) {metric <0-
16777214>|metric-type (1|2)|route-map WORD|tag <0-4294967295>}
no redistribute (kernel|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.

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

None

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ipv6 ospf Tag1
(config-router)#
```

© 2026 IP Infusion Inc. Proprietary

show debugging ipv6 ospf

Use this command to display the OSPFv3 debugging options.

Command Syntax

```
show debugging ipv6 ospf
```

Parameters

None

Command Mode

Execution mode and Privileged execution 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
```

© 2026 IP Infusion Inc. Proprietary

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 execution mode and Execution 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
```

The following table explains the fields for each ospf entry.

Table 91. 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.

Table 91. show ipv6 ospf output details (continued)

Field	Description
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.
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.

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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 execution mode and Execution 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

The following table explains the fields for each database entry.

Table 92. 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.

Table 92. OSPFv3 database fields (continued)

Field	Description
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.
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.

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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 execution mode and Execution 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
```

© 2026 IP Infusion Inc. Proprietary

```
    LS-Ack received 0 sent 3, Discarded 0
```

The following table explains the fields for each ospf interface entry.

Table 93. 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: 1. Ethernet is Broadcast 2. Serial p2p non-broadcast 3. 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.

© 2026 IP Infusion Inc. Proprietary

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 execution mode and Execution 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
```

© 2026 IP Infusion Inc. Proprietary

```
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
```

The following table explains the fields for each ospf neighbor entry.

Table 94. 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.
Dead timer due in	The countdown timer for marking neighbor connections dead. In this

Table 94. show ipv6 ospf neighbor output details (continued)

Field	Description
	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.

© 2026 IP Infusion Inc. Proprietary

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 (page 2600) 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 execution mode and Execution 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
```

© 2026 IP Infusion Inc. Proprietary

```
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
```

The following explains the fields for each ospf route entry.

Table 95. 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 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 execution mode and Execution 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#
```

The table explains the fields for each ospf virtual-links entry.

Table 96. 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.

© 2026 IP Infusion Inc. Proprietary

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 execution mode and Execution 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
```

The following table explains the fields for each ospfv3 topology entry.

Table 97. show ipv6 ospfv3 topology output details

Field	Description
OSPFv3 path to Area	Area ID in IPv4 format.
Router ID	ID in IPv4 format,
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.

© 2026 IP Infusion Inc. Proprietary

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 execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Usage

```
# show ipv6 route fast-reroute
```

© 2026 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

Execution modeand Privileged execution 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
```

The table explains the fields.

Table 98. 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
```

© 2026 IP Infusion Inc. Proprietary

snmp restart ospf6

Use this command restart SNMP in OSPFv3

Command Syntax

```
snmp restart ospf6
```

Parameter

None

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#snmp restart ospf6
```

© 2026 IP Infusion Inc. Proprietary

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

Zero

Command Mode

Router mode

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

Minimum delay: 500 milliseconds

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
```

© 2026 IP Infusion Inc. Proprietary

OSPF VPN Commands

This section provides an alphabetized reference of the OSPF VPN commands. It includes the following commands:

capability vrf-lite 2647
domain-id 2648
router ospf vrf 2650

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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)#
```

© 2026 IP Infusion Inc. Proprietary

CSPF-TE Commands

This section is a reference for Constrained Shortest Path First Traffic Engineering (CSPF-TE) commands:

cspf default-retry-interval 2652
cspf disable-better-protection 2653
cspf enable-better-protection 2654
cspf enable-resource-constraint-crankback 2655
cspf tie-break 2656

© 2026 IP Infusion Inc. Proprietary

cspf default-retry-interval

Use this command to set the default route computation retry interval for an OSPFv2 instance. This value is used for route recomputation (in the case of computation failures) and no retry interval is specified for a given LSP.

Use the `no` parameter with this command to unset default route computation retry interval.

Command Syntax

```
cspf default-retry-interval <1-3600>
no cspf default-retry-interval
```

Parameter

<1-3600>
The retry interval in seconds. The default is 10 seconds.

Default

None

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#cspf default-retry-interval 720
```

© 2026 IP Infusion Inc. Proprietary

cspf disable-better-protection

Use this command to disable better protection (Default).

Command Syntax

```
cspf disable-better-protection
```

Parameter

None

Default

None

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#cspf disable-better-protection
```

© 2026 IP Infusion Inc. Proprietary

cspf enable-better-protection

Use this command to enable better protection.

Command Syntax

```
cspf enable-better-protection
```

Parameter

None

Default

None

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#cspf enable-better-protection
```

© 2026 IP Infusion Inc. Proprietary

cspf enable-resource-constraint-crankback

Use this command to enable resource constraint checking for crank back when the selected path does not meet the requirements.

Use cspf disable-resource-constraint-crankback to disable resource constraint checking for crank back when the selected path does not meet the requirements.

Command Syntax

```
cspf enable-resource-constraint-crankback
cspf disable-resource-constraint-crankback
```

Parameter

None

Default

cspf disable-resource-constraint-crankback

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)# cspf enable-resource-constraint-crankback
```

© 2026 IP Infusion Inc. Proprietary

cspf tie-break

Use this command to set the tie-break method for an OSPFv2 instance. This selects a link (during route computation) when more than one candidate link satisfies all the route constraints and the associated cost and hop limit link attributes are equal.

Use the `no` parameter with this command to unset tie-break method.

Command Syntax

```
cspf tie-break (random|least-fill|most-fill)
no cspf tie-break
```

Parameters

`random`
Pick any path at random. This method places an equal number of LSPs on each link without taking into account the available bandwidth ratio. This is the default tie-break method.

`least-fill`
Pick the path with the largest minimum available bandwidth ratio. This method equalizes the reservation on each link.

`most-fill`
Pick the path with the smallest minimum available bandwidth ratio. This method uses one link until it is completely full and then uses the next link.

Default

By default, the tie-break method is set to `random`.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#cspf tie-break least-fill
```

© 2026 IP Infusion Inc. Proprietary

ROUTING INFORMATION PROTOCOL CONFIGURATION

Router Information Protocol Configuration 2658
Enable RIP 2658
Specify RIP Version 2661
Authentication with a Single Key 2664
Text Authentication with Multiple Keys 2669
MD5 Authentication with Multiple Keys 2677
RIPV2 VRF Configuration 2683
RIPng 2686
Topology 2686
Configuration 2686
Validation 2687

© 2026 IP Infusion Inc. Proprietary

Router Information Protocol Configuration

This section 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 183. Enable RIP Topology

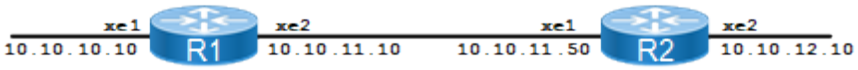

R1

`#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

© 2026 IP Infusion Inc. Proprietary

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 184. RIP Version Topology

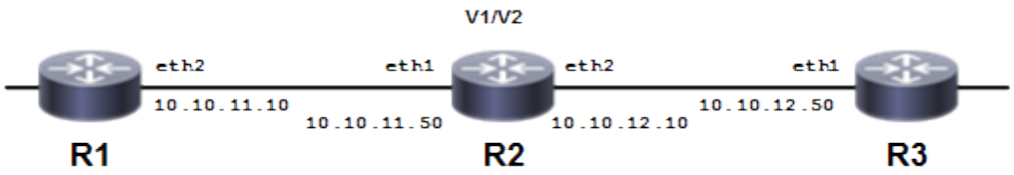

R2

`#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 185. Single-key Topology

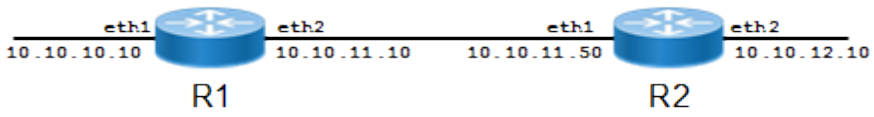

R1

`#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.

© 2026 IP Infusion Inc. Proprietary

Figure 186. Multiple-key Topology

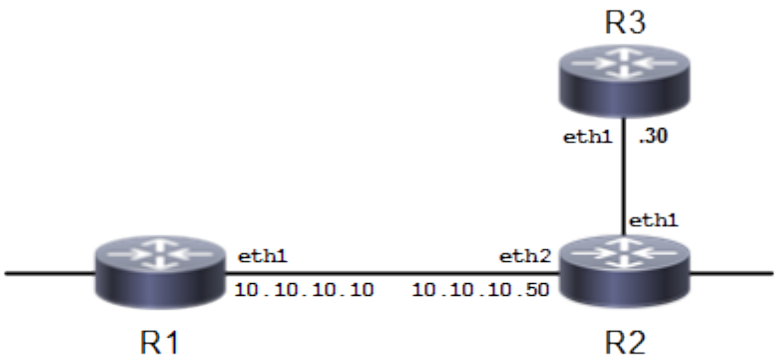

R1

`#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

© 2026 IP Infusion Inc. Proprietary

`(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
```

© 2026 IP Infusion Inc. Proprietary

```
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 187. MD5 Multiple-key Topology

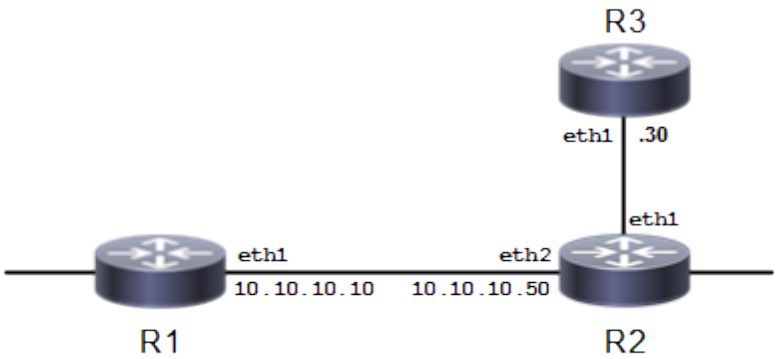

R1

`#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

© 2026 IP Infusion Inc. Proprietary

	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

© 2026 IP Infusion Inc. Proprietary

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
```

RIPV2 VRF Configuration

RIPV2 can be configured along with VRF between two nodes and also between PE-CE nodes.

Topology

Figure 188. RIPV2 VRF Topology

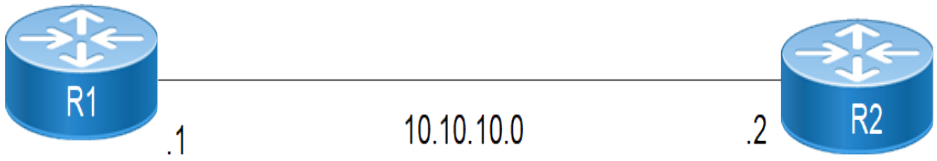

RIPV2 VRF Configuration

This document captures requirements to use RIPv2 with VRF between two nodes.

R1

`#configure terminal`	Enter configuration mode.
`(config)#ip vrf ripv2`	Configure ip vrf
`(config-vrf)#rd 1:1`	Configure rd
`(config-vrf)#route-target both 1:100`	Configure rt value
`(config-vrf)#int xe48`	Enter in to interface
`(config-if)#ip vrf forwarding ripv2`	Enable ip vrf forwarding
`(config-if)#ip address 10.10.10.1/24`	Configure ip address
`(config-if)#router rip`	Enter in to router rip
`(config-router)#address-family ipv4 vrf ripv2`	Address family ipv4 with vrf name
`(config-router-af)#network 10.10.10.0/24`	Configure network command
`(config-router-af)#redistribute static`	Configure Redistribute static
`(config-router-af)#redistribute connected`	Configure Redistribute connected
`(config-router-af)#commit`	Commit the transactions
`(config)#ip route vrf ripv2 191.1.1.10/32 xe48`	Configure static route with vrf
`(config)#commit`	Commit the transaction

R2

`#configure terminal`	Enter configuration mode.
`(config)#ip vrf ripv2`	Configure ip vrf
`(config-vrf)#rd 1:1`	Configure rd
`(config-vrf)#route-target both 1:100`	Configure rt value
`(config-vrf)#int xe10`	Enter in to interface
`(config-if)#ip vrf forwarding ripv2`	Enable ip vrf forwarding
`(config-if)#ip address 10.10.10.2/24`	Configure ip address
`(config-if)#router rip`	Enter in to router rip

© 2026 IP Infusion Inc. Proprietary

`(config-router)#address-family ipv4 vrf ripv2`	Address family ipv4 with vrf name
`(config-router-af)#network 10.10.10.0/24`	Configure network command
`(config-router-af)#redistribute static`	Configure Redistribute static
`(config-router-af)#redistribute connected`	Configure Redistribute connected
`(config-router-af)#commit`	Commit the transactions

Validation

R1

```
#sh ip rip database vrf ripv2

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                  xe48
S  191.1.1.10/32                           1                  xe48
```

R2

```
#sh ip rip database vrf ripv2

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                  xe10
R  191.1.1.10/32      10.10.10.1           2 10.10.10.1       xe10   02:48
```

© 2026 IP Infusion Inc. Proprietary

RIPng

This section contains a basic RIPng configuration example.

For details about the commands used in these examples, see the *Routing Information Protocol Command Reference (page 2689)*.

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 189. RIPng Topology

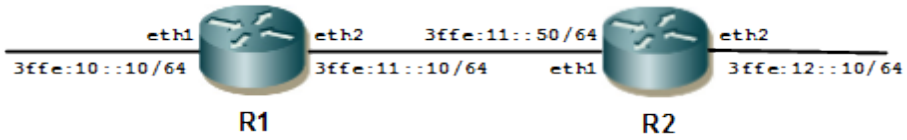

Configuration

R1

`#configure terminal`	Enter Configure mode.
`(config)#interface eth1`	Specify the interface (`eth1`) to configure, and 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`	Specify the interface (`eth2`) to configure, and 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.
`(config)# commit`	Commit the transaction.

R2

`#configure terminal`	Enter Configure mode.
`(config)#interface eth1`	Specify the interface (`eth1`) to configure, and enter Interface mode.
`(config-if)#ipv6 router rip`	Enable RIPng routing on interface `eth1`.

© 2026 IP Infusion Inc. Proprietary

`(config-if)#exit`	Exit Interface mode, and enter Configure mode.
`(config)#interface eth2`	Specify the interface (`eth2`) to configure, and 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.
`(config)# commit`	Commit the transaction.

Validation

R1

```
R1#show ipv6 rip

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
Rc 3ffe:10::/64                ::                          eth1     1   0
Rc 3ffe:11::/64                ::                          eth2     1   0
R  3ffe:12::/64                fe80::aa2b:b5ff:fe1c:c561   eth2     2   0  02:40

R1#show ipv6 rip interface
lo.management is up, line protocol is up
  RIPng is not enabled on this interface
eth1 is up, line protocol is up
  Routing Protocol: RIPng
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IPv6 interface address:
      3ffe:11::10/64
      fe80::aa2b:b5ff:fe2f:41cb/64
eth2 is up, line protocol is up
  Routing Protocol: RIPng
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IPv6 interface address:
      3ffe:10::10/64
      fe80::aa2b:b5ff:fe2f:41cb/64

R1#show ipv6 protocols rip

Routing Protocol is "ripng"

  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 redistribute metric is 1
  Redistributing:
  Interface
    xe48
    ce49
  Routing for Networks:

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, 4d19h49m
C      3ffe:10::/64 via ::, eth1, 00:10:53
C      3ffe:11::/64 via ::, eth2, 00:10:08
R      3ffe:12::/64 [120/2] via fe80::aa2b:b5ff:fe1c:c561, eth2, 00:04:26
C      fe80::/64 via ::, eth2, 00:54:20
```

R2

```
R2#show ipv6 rip

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
R  3ffe:10::/64                fe80::aa2b:b5ff:fe2f:41cb   eth1     2   0  02:36
Rc 3ffe:11::/64                ::                          eth1     1   0
Rc 3ffe:12::/64                ::                          eth2     1   0

R2#show ipv6 rip interface
eth1 is up, line protocol is up
  Routing Protocol: RIPng
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IPv6 interface address:
      3ffe:11::50/64
      fe80::aa2b:b5ff:fe1c:c561/64
eth2 is up, line protocol is up
  Routing Protocol: RIPng
    Passive interface: Disabled
    Split horizon: Enabled with Poisoned Reversed
    IPv6 interface address:
      3ffe:12::10/64
      fe80::aa2b:b5ff:fe1c:c561/64
```

ROUTING INFORMATION PROTOCOL COMMAND REFERENCE

Routing Information Protocol Commands 2691
accept-lifetime 2693
cisco-metric-behavior 2695
clear ip rip route 2696
clear ip rip route vrf NAME 2698
clear ip rip statistics 2699
debug rip 2700
default-information originate 2702
default-metric 2703
distance 2704
distribute-list 2705
ip rip authentication key-chain 2706
ip rip authentication mode 2707
ip rip authentication string 2708
ip rip receive-packet 2709
ip rip receive version 2710
ip rip send-packet 2711
ip rip send version 2712
ip rip split-horizon 2713
key chain 2714
key-id 2715
key-string 2716
key-string encrypted 2717
maximum-prefix 2718
neighbor 2719
network 2720
offset-list 2721
passive-interface 2722
recv-buffer-size 2723
redistribute 2724
route 2726
router rip 2727
send-lifetime 2728
show debugging rip 2730
show ip protocols rip 2731

© 2026 IP Infusion Inc. Proprietary

show ip rip 2733
show ip rip interface 2735
show ip rip statistics 2737
snmp restart rip 2739
timers basic 2740
version 2741
RIPng Commands 2742
aggregate-address 2743
cisco-metric-behavior 2744
clear ipv6 rip route 2745
debug ipv6 rip 2746
default-information originate 2747
default-metric 2748
distance 2749
distribute-list 2750
ipv6 rip metric-offset 2751
ipv6 rip split-horizon 2752
ipv6 router rip 2753
neighbor 2754
offset-list 2755
passive-interface 2756
recv-buffer-size 2757
redistribute 2758
route 2760
route-map 2761
router ipv6 rip 2762
show debugging ipv6 rip 2763
show ipv6 protocols rip 2764
show ipv6 rip 2765
show ipv6 rip interface 2766
timers basic 2767
Routing Information Protocol VPN Commands 2768
show ip rip interface vrf 2769
show ip rip vrf 2771
show ip vrf 2773
Routing Information Protocol Authentication 2774
Single Key Authentication 2774
Multiple Keys Authentication 2774

© 2026 IP Infusion Inc. Proprietary

Routing Information Protocol Commands

This section provides an alphabetized reference for each of the Routing Information Protocol (RIP) commands, which support IPv4. It includes the following commands:

accept-lifetime 2693
cisco-metric-behavior 2695
clear ip rip route 2696
clear ip rip route vrf NAME 2698
clear ip rip statistics 2699
debug rip 2700
default-information originate 2702
default-metric 2703
distance 2704
distribute-list 2705
ip rip authentication key-chain 2706
ip rip authentication mode 2707
ip rip authentication string 2708
ip rip receive-packet 2709
ip rip receive version 2710
ip rip send-packet 2711
ip rip send version 2712
ip rip split-horizon 2713
key chain 2714
key-id 2715
key-string 2716
key-string encrypted 2717
maximum-prefix 2718
neighbor 2719
network 2720
offset-list 2721
passive-interface 2722
recv-buffer-size 2723
redistribute 2724
route 2726
router rip 2727
send-lifetime 2728
show debugging rip 2730
show ip protocols rip 2731
show ip rip 2733

© 2026 IP Infusion Inc. Proprietary

show ip rip interface 2735
show ip rip statistics 2737
snmp restart rip 2739
timers basic 2740
version 2741

© 2026 IP Infusion Inc. Proprietary

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 A, *Routing Information Protocol Authentication* for information on how this command is related to the other authentication commands.

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
```

Parameter

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 time when accept-lifetime expires 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 expire.

duration
Specify the duration of the key in seconds <1-2147483646>.

infinite
Specify the end time to never expire.

Default

None

© 2026 IP Infusion Inc. Proprietary

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 `key-id 1` on the key chain named `mychain`.

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key-id 1
(config-keychain-key)#accept-lifetime 03:03:01 Dec 30 2004 04:04:02 Oct 06
2006
(config)#key chain mychain
(config-keychain)#key-id 1
(config-keychain-key)#no accept-lifetime
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

Examples

```
#clear ip rip route 10.0.0.0/8
#clear ip rip route ospf
```

© 2026 IP Infusion Inc. Proprietary

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 *
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

```
#configure terminal
(config)#debug rip events
(config)#debug rip packet send detail
(config)#debug rip nsm
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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>
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

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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.

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
```

© 2026 IP Infusion Inc. Proprietary

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.

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
```

© 2026 IP Infusion Inc. Proprietary

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.

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
ino 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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 A, Routing Information Protocol Authentication (page 2774) for information on how this command is related to the other authentication commands.

Command Syntax

```
key chain WORD
no key chain WORD
```

Parameters

WORD
Specify the name of the key chain to manage.

Default

None

Command Mode

Configure mode and Key-chain 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 `key-id 10` in the same mode.

```
(config-keychain)#key chain mykeychain3
(config-keychain)#key-id 10
(config-keychain-key)#
```

© 2026 IP Infusion Inc. Proprietary

key-id

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 A, *Routing Information Protocol Authentication* for information on how this command is related to the other authentication commands.

Command Syntax

```
key-id <0-2147483647>
no key-id <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-id 1` in the key chain named `mychain` is set to `prime`:

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key-id 1
(config-keychain-key)#key-string prime
(config-keychain)#key-id 1
(config-keychain-key)#no key-string
```

© 2026 IP Infusion Inc. Proprietary

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.

See Appendix A, *Routing Information Protocol Authentication* for information on how this command is related to the other authentication commands.

Command Syntax

```
key-string WORD
no key-string
```

Parameters

WORD
Specify a string of characters to be used as a password by the key.

Default

Disabled

Command Mode

Keychain mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

In the following example, the password for `key-id 1` in the key chain named `mychain` is set to prime:

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key-id 1
(config-keychain-key)#key-string prime

(config-keychain)#key-id 1
(config-keychain-key)#no key-string
```

© 2026 IP Infusion Inc. Proprietary

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

Key-chain mode and Key-chain key mode

Applicability

This command was introduced in OcNOS version 4.1.

Example

In the following example, the encrypted password for key-id 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-id 1
(config-keychain-key)#key-string encrypted 0xd6c50b442de47f70
(config-keychain)#key-id 1
(config-keychain-key)#no key-string
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 (kernel|connected|static|ospf|isis|bgp)
redistribute (kernel|connected|static|ospf|isis|bgp) metric <0-16>
redistribute (kernel|connected|static|ospf|isis|bgp) route-map WORD
redistribute (kernel|connected|static|ospf|isis|bgp) metric <0-16> route-map WORD
no redistribute (kernel|connected|static|ospf|isis|bgp)
no redistribute (kernel|connected|static|ospf|isis|bgp) metric <0-16>
no redistribute (kernel|connected|static|ospf|isis|bgp) route-map WORD
no redistribute (kernel|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.

© 2026 IP Infusion Inc. Proprietary

Examples

```
#configure terminal
(config)#router rip
(config-router)#redistribute connected

#configure terminal
(config)#router rip
(config-router)#redistribute connected route-map myroutemap
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 time when send-lifetime expires 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 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

© 2026 IP Infusion Inc. Proprietary

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example shows the setting of `send-lifetime` for `key-id 1` on the key chain named `mychain`:

```
#configure terminal
(config)#key chain mychain
(config-keychain)#key-id 1
(config-keychain-key)#send-lifetime 03:03:01 Jan 03 2004 04:04:02 Dec 06 2006
```

© 2026 IP Infusion Inc. Proprietary

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#
```

© 2026 IP Infusion Inc. Proprietary

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)
#
```

Table 99 Explains the show command output details.

Table 99. 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.

© 2026 IP Infusion Inc. Proprietary

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 100 shows the status codes displayed at the start of a route entry.

Table 100. 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.

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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
```

Table 101 Explains the show command output details.

Table 101. 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

© 2026 IP Infusion Inc. Proprietary

	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.

© 2026 IP Infusion Inc. Proprietary

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
```

Table 102 Explains the show command output details.

Table 102. 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.

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

RIPng Commands

This section 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 2743
cisco-metric-behavior 2744
clear ipv6 rip route 2745
debug ipv6 rip 2746
default-information originate 2747
default-metric 2748
distance 2749
distribute-list 2750
ipv6 rip metric-offset 2751
ipv6 rip split-horizon 2752
ipv6 router rip 2753
neighbor 2754
offset-list 2755
passive-interface 2756
recv-buffer-size 2757
redistribute 2758
route 2760
route-map 2761
router ipv6 rip 2762
show debugging ipv6 rip 2763
show ipv6 protocols rip 2764
show ipv6 rip 2765
show ipv6 rip interface 2766
timers basic 2767

© 2026 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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear ipv6 rip route isis
#clear ipv6 rip route 3ffe:ffff::/16
```

© 2026 IP Infusion Inc. Proprietary

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 execution 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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 2703)).

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 2704)).

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 me tric 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 2713)).

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 2719)).

Command Syntax

```
neighbor X:X::X:X IFNAME
no neighbor X:X::X:X IFNAME
```

Parameters

X:X::X:X

Specify a link-local address of the peer 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 fe80::5054:ff:fece:30cd eth1
(config-router)#end
#show running-config ripng
!
router ipv6 rip
neighbor fe80::5054:ff:fece:30cd eth1
!
#

#show ipv6 interface brief
Interface            IPv6-Address                                    Admin-Status
eth0                 fe80::5054:ff:fe4e:22d3                         [up/up]
eth1                 2001:2001::1:2
                         fe80::5054:ff:fe5c:6a4f                         [up/up]
eth2                 fe80::5054:ff:fe16:8339                         [up/up]
eth3                 fe80::5054:ff:feb9:bc7c                         [up/up]
lo                   ::1                                             [up/up]
```

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 (page 2721)).

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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 (kernel|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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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)#
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution 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

Execution mode and Privileged execution 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
#
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution 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
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution 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
```

© 2026 IP Infusion Inc. Proprietary

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 (page 2740)).

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
```

© 2026 IP Infusion Inc. Proprietary

Routing Information Protocol VPN Commands

This section 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 MPLS/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 section contains the following commands:

show ip rip interface vrf 2769
show ip rip vrf 2771
show ip vrf 2773

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution 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
```

The following Explains the show command output details.

Table 103. 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.

© 2026 IP Infusion Inc. Proprietary

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.

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution 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
```

The following table shows the status codes displayed at the start of a route entry.

Table 104. 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.

© 2026 IP Infusion Inc. Proprietary

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.

The following table Explains the show command output details.

Table 105. 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.

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show ip vrf VRF_A
VRF VRF_A; (table=1)
```

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-id 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 WORD
```

where `WORD` is a string of characters to be used as a password by the key.

4. Set key management options

© 2026 IP Infusion Inc. Proprietary

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-id 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-id 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 08 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
```

© 2026 IP Infusion Inc. Proprietary

TWO-WAY ACTIVE MEASUREMENT PROTOCOL CONFIGURATION

Two-way Active Measurement Protocol 2777
Overview 2777
Topology 2778
Configuration 2778
End-to-End Measurement 2782
TWAMP for Client and Server with SNMP 2787
Overview 2787
TWAMP Client 2787
TWAMP Server 2787
SNMP for TWAMP 2788
Benefits 2788
Server Configuration 2789
Configuration to Trigger TWAMP Traps 2790
Glossary 2800
TWAMP for VPN Inter-Access Interface 2801
Overview 2801
Feature Characteristics 2801
Benefits 2801
Prerequisites 2801
Limitations 2801
Topology 2802
Configurations 2802
CLI Commands 2809
Glossary 2838

© 2026 IP Infusion Inc. Proprietary

Two-way Active Measurement Protocol

Overview

Two-way Active Measurement Protocol (TWAMP) is an open protocol for measuring network performance between any two devices.

In OcNOS, TWAMP measurement is configured under the interface mode using using the “delay-measurement dynamic twamp” command. When the command is configured, the measurement is started on the corresponding interface and using the configuration data in the delay-measurement command. The measurement is stopped by unconfiguring the “delay-measurement” command on the interface.

The TWAMP hardware profile filter should be enabled so that the timestamps used in the TWAMP protocol calculations are obtained accurately from the DNX hardware.

After the measurement is done, the result can be checked by the user using the OcNOS cli or netconf. The result is also advertised to IGP TE. There are 2 modes of advertisement, periodic and accelerated. Both periodic and accelerated advertisements can be enabled in parallel.

Periodic Advertisement

Periodic advertisement is enabled by default. In this case, the link delay metrics are calculated at a periodic “computation” interval. The calculated minimum delay for a link is compared with the minimum delay value advertised previously:

- The delay metrics for the link are advertised only if the variation in values is beyond any of the configured periodic thresholds.
- If the variation in values is within configured periodic thresholds, the delay metrics for the link are not advertised.

Accelerated Advertisement

Accelerated advertisement is disabled by default. In this case, the link delay metrics are calculated after receiving every burst of data. The calculated minimum delay for a link is compared with the minimum delay value advertised previously:

- The delay metrics for the link are advertised only if the variation in values is beyond any of the configured accelerated thresholds.
- If the variation in values is within configured accelerated thresholds, the delay metrics for the link are not advertised.

Limitation and Assumptions

- To verify delay, loss and bandwidth measurement with IGP_TE, the pre-requisites are that the ISIS and OSPF sessions should be up and running. (Check OSPF and ISIS configuration in corresponding sections).
- The a-bit-threshold configures the Anomalous (A) bit. The A bit is set when the measured value of this parameter exceeds its configured maximum threshold. The A bit is cleared when the measured value falls below its configured reuse threshold.
- The “test-session-name” command is no longer supported and should not be used. The “delay-measurement” command under interface mode should be used instead.

© 2026 IP Infusion Inc. Proprietary

- On 9600-56DX boards, if TWAMP and VXLAN are used simultaneously, the hardware profile for NVO VXLAN should be configured last.

Topology

Figure 190. TWAMP topology

Configuration

Configuration TWAMP Measurement on ROUTER

Configure the router R1 to send TWAMP measurement packets on interface xe1 towards the reflector on R2.

R1

`#configure terminal`	Enter configure mode.
`(config)# hardware-profile filter twamp-ipv4 enable`	Enable the hardware profile filter for TWAMP ipv4 so that HW timestamps can be retrieved in the DNX hardware for IPv4.
`(config)# hardware-profile filter twamp-ipv6 enable`	Enable the hardware profile filter for TWAMP ipv6 so that HW timestamps can be retrieved in the DNX hardware for IPv6.
`(config)#twamp-light control`	Enter to Twamp-Light control mode.
`(config-twamp-light-ctrl)#control-admin-state enable`	Enable the Twamp-light sender framework.
`(config-twamp-light-ctrl)#exit`	Exit the Twamp-light control mode.
`(config)#interface xe1`	Enter interface mode.
`(config-if)#ip address 23.1.1.1/24`	Configure ipv4 address.
`(config-if)#ipv6 address 222::1/64`	Configure ipv6 address.
`(config-if)# delay-measurement dynamic twamp reflector-ip 23.1.1.2 reflector-port 1235 sender-ip 23.1.1.1 sender-port 1234 dscp 56`	Configure the TWAMP delay measurement.
`(config-if)#commit`	Commit the candidate configuration to running configuration.

Configure the router R2 to receive and reflect TWAMP measurement packets on interface xe2.

R2

`#configure terminal`	Enter configure mode.
`(config)# hardware-profile filter twamp-ipv4`	Enable the hardware profile filter for TWAMP ipv4 so

`enable`	that HW timestamps can be retrieved in the DNX hardware for IPv4.
`(config)# hardware-profile filter twamp-ipv6 enable`	Enable the hardware profile filter for TWAMP ipv4 so that HW timestamps can be retrieved in the DNX hardware for IPv6.
`(config)#interface xe2`	Enter interface mode.
`(config-if)#ip address 23.1.1.2/24`	Configure ipv4 address.
`(config-if)#ipv6 address 222::2/64`	Configure ipv6 address.
`(config-if)#exit`	Exit interface mode
`(config)#twamp-light reflector`	Enter to Twamp-Light reflector mode
`(config-twamp-light-ref)#reflector-admin-state enable`	Enable the Twamp-light reflector framework
`(config-twamp-light-ref)#reflector-name r1 reflector-ip 23.1.1.2 reflector-port 1235 sender-ip 23.1.1.1 sender-port 1234`	Create a reflector instance
`(config-twamp-light-ref)#commit`	Commit the candidate configuration to the running configuration
`(config-twamp-light-ref)#exit`	Exit Twamp-Light reflector mode

Configuring TWAMP Delay Profiles

The delay profile “interfaces” configures the way the TWAMP packets are sent (burst, interval…) and the way the data is advertised after the measurement data is collected

R1

`#configure terminal`	Enter configure mode.
`(config)# delay-profile interfaces`	Enter in to delay profile mode
`(config-dp-int)#burst interval 3000`	To config the burst interval for TWAMP packets in milliseconds(default value is 3000 milliseconds).
`(config-dp-int)#burst count 10`	To set the number of packets to be sent at each burst interval (default=10).
`(config-dp-int)#interval 30`	To set the computation interval for TWAMP packets in seconds. (Default=30 sec).
`(config-dp-int)#interval 120`	To set the computation interval of TWAMP packets for advertisement periodic mode.
`(config-dp-int)# advertisement periodic threshold 10`	Config threshold value inside advertisement periodic mode for TWAMP packets.
`(config-dp-int)# advertisement periodic minimum-change 1000`	Config minimum change value inside advertisement periodic mode for TWAMP packets.
`(config-dp-int)# advertisement accelerated`	Enable advertisement accelerated in delay profile mode.
`(config-dp-int)# advertisement accelerated threshold 20`	Config advertisement accelerated minimum change value inside delay profile mode for TWAMP packets.

© 2026 IP Infusion Inc. Proprietary

`(config-dp-int)# advertisement accelerated minimum-change 2000`	Config advertisement accelerated minimum change value inside delay profile mode for TWAMP packets.
`(config--dp-int)#commit`	Commit the candidate configuration to running configuration

Configuring Loss Measurement

Loss measurement is disabled by default and should be configured separately after the delay-measurement has been configured.

R1

`#configure terminal`	Enter configure mode
`(config)#interface xe1`	Enter interface mode
`(config-if)#delay-measurement dynamic twamp reflector-ip 23.1.1.2 reflector-port 1235 sender-ip 23.1.1.1 sender port 1234 dscp 56`	Configure delay measurement
`(config-if)#loss-measurement dynamic`	Configure loss measurement after delay measurement is already configured
`(config-if)#commit`	Commit the candidate configuration to running configuration
`(config-if)#exit`	Exit interface mode

Un-Configuring/Stopping TWAMP Measurement

R1

`#configure terminal`	Enter configure mode
`(config)# hardware-profile filter twamp-ipv4 disable`	Disable hardware profile filter for TWAMP ipv4.
`(config)# hardware-profile filter twamp-ipv6 disable`	Disable hardware profile filter for TWAMP ipv6.
`(config)#interface xe1`	Enter interface mode
`(config-if)# no delay-measurement dynamic twamp reflector-ip 23.1.1.2`	Unconfigure delay measurement.
`(config-if)#no loss-measurement dynamic`	Unconfigure loss measurement.
`(config-if)#commit`	Commit the candidate configuration to running configuration.
`(config-if)#exit`	Exit interface mode.

Un-Configuring TWAMP Delay profile

R1

`#configure terminal`	Enter configure mode
`(config)# delay-profile interfaces`	Enter in to delay profile mode.

(config-dp-int)#no burst interval	Un-Configure burst interval value inside delay profile mode.
(config-dp-int)#no burst count	Un-Configure burst count value inside delay profile mode.
(config-dp-int)#no interval	Un-Configure interval value inside delay profile mode.
(config-dp-int)#no sender-port	Un-Configure sender port value inside delay profile mode.
(config-dp-int)#no advertisement periodic	Disable advertisement periodic mode.
config-dp-int)#no interval	Un-Configure interval value inside for advertisement periodic mode.
(config-dp-int)#no advertisement periodic threshold	Un-Configure threshold value inside advertisement periodic mode
(config-dp-int)# no advertisement periodic minimum-change	Un-Configure minimum change value inside advertisement periodic mode
(config-dp-int)#no advertisement accelerated	Un-Configure accelerated inside advertisement mode
(config-dp-int)#no advertisement accelerated minimum-change	Un-Configure advertisement accelerated minimum change value inside delay profile mode
(config-dp-int)# no advertisement accelerated threshold	Un-Config advertisement accelerated threshold value inside delay profile mode
(config-dp-if)#commit	Commit the candidate configuration to running configuration

Clearing TWAMP Statistics

R1

Router1# clear twamp measurement interface xe1	This command will restart the measurement on the specified interface, resetting the collected stats.
Router1# clear twamp measurement interface all	This command will restart the measurement on all the interfaces, resetting all the collected stats. If there are any non-applied configuration changes to the delay profile, they will be applied when this command is issued.

Validation

Configuration Validation on the Sender Side

```
#show running-config twamp
twamp-light control
 control-admin-state enable
#show running-config interface xe1 | inc delay-measurement
  delay-measurement
 dynamic twamp
 reflector-ip 23.1.1.2
 reflector-port 1235
 sender-ip 23.1.1.1
 sender-port 1234
 dscp 56
```

© 2026 IP Infusion Inc. Proprietary

Configuration validation on reflector side

```
#show running config twamp
twamp-light reflector
 reflector-admin-state enable
 reflector-name r1
 reflector-ip 23.1.1.2
 reflector-port 1234
```

Validation of TWAMP statistics

```
#show twamp-statistics interfaces xe1
Interface name    : xe1
Sender IP         : 23.1.1.1
Reflector IP      : 23.1.1.2
Reflector port    : 1234
DSCP value        : 56
Last Advertised stats:
  Time: 2019-02-14 10:19:38
  Average delay          : 2
  Minimum delay          : 2
  Maximum delay          : 2
  Average delay variation: 0
  Minimum delay variation: 0
  Maximum delay variation: 0
  Packets sent           : 100
  Packets received       : 100
  Packets timeout        : 0
  Packet Loss: 0
Last Calculated stats:
  Time: 2019-02-14 10:23:39
  Average delay          : 2
  Minimum delay          : 2
  Maximum delay          : 2
  Average delay variation: 0
  Minimum delay variation: 0
  Maximum delay variation: 0
  Packets sent           : 100
  Packets received       : 100
  Packets timeout        : 0
  Packet Loss : 0
```

End-to-End Measurement

When a reflector is multiple hops away from the sender and there are multiple paths available to the reflector, End-to-End measurement ensures that the TWAMP packets should go out of only the interface with the configured Sender IP/Reflector IP on sender and reflector side respectively.

© 2026 IP Infusion Inc. Proprietary

Topology

Figure 191. TWAMP topology

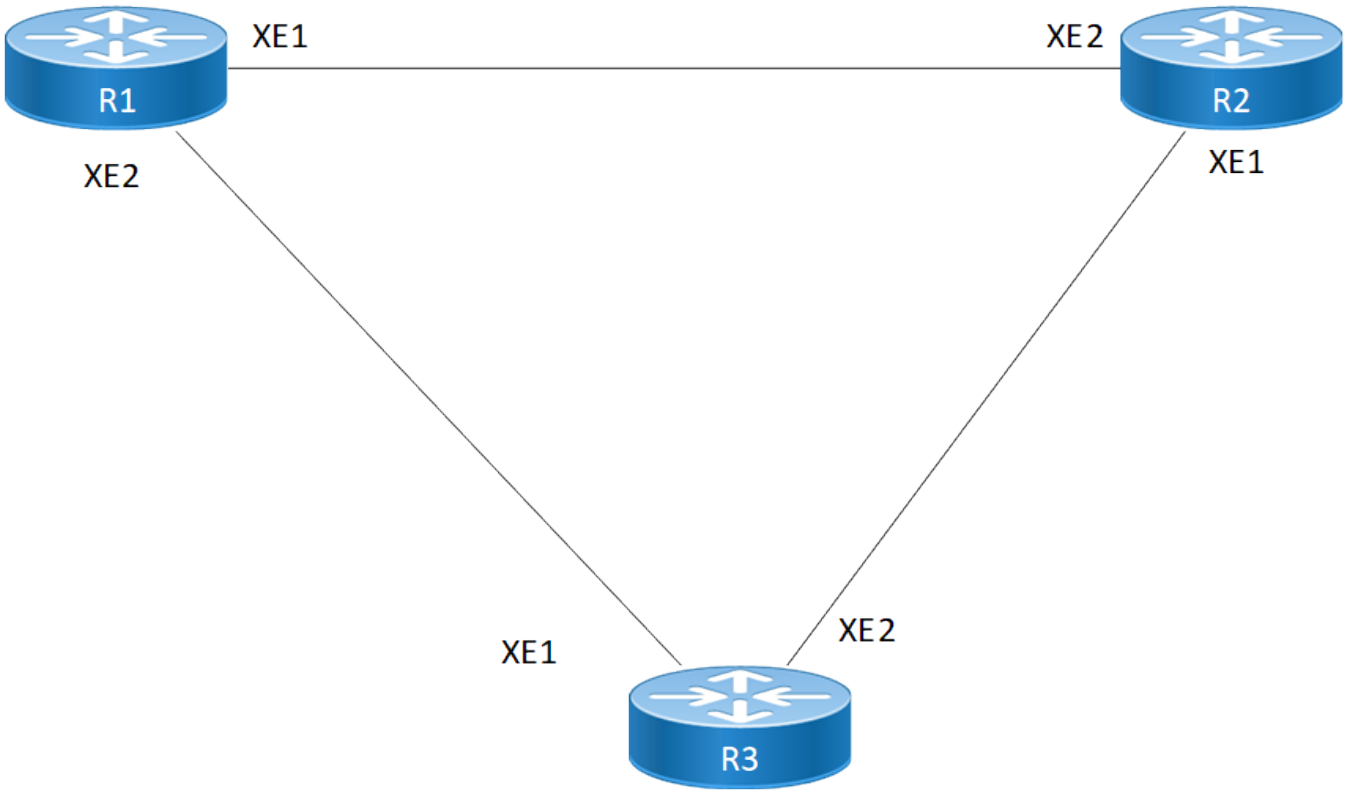

Configuration

R1 - Sender

`#configure terminal`	Enter configure mode.
`(config)# hardware-profile filter twamp-ipv4 enable`	Enable the hardware profile filter for TWAMP ipv4 so that HW timestamps can be retrieved in the DNX hardware for IPv4.
`(config)#twamp-light control`	Enter the Twamp-Light control mode.
`(config-twamp-light-ctrl)#control-admin-state enable`	Enable the Twamp-light sender framework.
`(config-twamp-light-ctrl)#exit`	Exit the Twamp-light Control mode.
`(config)#interface xe1`	Enter interface mode.
`(config-if)#ip address 10.10.10.1/24`	Configure ipv4 address.
`(config-if)#exit`	Exit interface mode.
`(config)#interface xe2`	Enter interface mode.
`(config-if)#ip address 20.20.20.1/24`	Configure ipv4 address.
`(config-if)# delay-measurement dynamic twamp sender-ip 20.20.20.1 reflector-ip 30.30.30.2 reflector-port 1234`	Configuring delay measurement. Since it is applied on interface xe2, it will force the packets to go out from

	xe2 even if others paths are available.
(config-if)#commit	Commit the candidate configuration to running configuration.
(config-if)#exit	Exit interface mode
(config)#ip route 30.30.30.0/24 20.20.20.2	Set static route to reflector.
(config)#commit	Commit the candidate configuration to running configuration.

R2 - Reflector

#configure terminal	Enter configure mode.
(config)# hardware-profile filter twamp-ipv4 enable	Enable the hardware profile filter for TWAMP ipv4 so that HW timestamps can be retrieved in the DNX hardware for IPv4.
(config)#interface xe1	Enter interface mode.
(config-if)#ip address 30.30.30.2/24	Configure ipv4 address.
(config)#interface xe2	Enter interface mode.
(config-if)#ip address 10.10.10.2/24	Configure ipv4 address.
(config-if)#exit	Exit interface mode.
(config)#twamp-light reflector	Enter the Twamp-Light reflector mode.
(config-twamp-light-ref)#reflector-admin-state enable	Enabling the Twamp-light reflector framework.
(config-twamp-light-ref)#reflector-name r1 reflector-ip 30.30.30.2 reflector-port 1234 sender-ip 20.20.20.1	Configuring the reflector instance with reflector ip and port number. Specifying explicitly the sender-ip in the reflector command will force the reflector to reply from the interface where the reflector-ip is configured.
(config-twamp-light-ref)#commit	Commit the candidate configuration to the running configuration.
(config-twamp-light-ref)#exit	Exit Twamp-light reflector framework mode.
(config-if)#ip route20.20.20.0/24 30.30.30.1	Set static route to reflector.
(config)#commit	Commit the candidate configuration to running configuration.

R3 - Intermediate

(config-twamp-light-ref)#reflector-admin-state enable	Enabling the Twamp-light reflector framework.
(config)#interface xe1	Enter xe1 interface mode
(config-if)#ip address 20.20.20.2/24	Configure ipv4 address
(config)#interface xe2	Enter xe2 interface mode
(config-if)#ip address 30.30.30.1/24	Configure ipv4 address
(config)#commit	Commit

© 2026 IP Infusion Inc. Proprietary

Validation

Validation On Sender Side

```
R1#show running-config twamp
twamp-light control
 control-admin-state enable

#show twamp-statistics interfaces xe2
Interface name    : xe2
Sender IP         : 20.20.20.1
Reflector IP      : 30.30.30.2
Reflector port    : 1234
DSCP value        : 0
HW Status         : -
Last Advertised stats:
  Time: 2019-03-07 10:45:54
  Average delay          : 628
  Minimum delay          : 558
  Maximum delay          : 697
  Average delay variation: 17
  Minimum delay variation: 8
  Maximum delay variation: 23
  Packets sent           : 50
  Packets received       : 50
  Packets timeout        : 0
  Packet Loss: Not Enabled
Last Calculated stats:
  Time: 2019-03-07 10:48:34
  Average delay          : 646
  Minimum delay          : 494
  Maximum delay          : 701
  Average delay variation: 8
  Minimum delay variation: 2
  Maximum delay variation: 21
  Packets sent           : 50
  Packets received       : 50
  Packets timeout        : 0
Packet Loss : Not Enabled

#show twamp-statistics
========================================
      TWAMP Test-Session Statistics
========================================
  Test-Session Name        : __internal_interface_10024
  Start Time               : 2019 Mar 07 10:43:13
  Elapsed time(milli sec)  : 348057
  Packets Sent             : 435
  Packets Received         : 435
  Packet Loss(%)           : 0.0000
  Round Trip Delay(usec)
      Minimum              : 607
      Maximum              : 3373
      Average              : 1216
  Forward Delay(usec)
      Minimum              : (*)
      Maximum              : (*)
      Average              : (*)
  Reverse Delay(usec)
      Minimum              : (*)
      Maximum              : (*)
      Average              : (*)
  Round Trip Delay Variation(usec)
      Minimum              : 144
      Maximum              : 852
      Average              : 243
  Forward Delay Variation(usec)
```

```
      Minimum              : (*)
      Maximum              : (*)
      Average              : (*)
  Reverse Delay Variation(usec)
      Minimum              : (*)
      Maximum              : (*)
      Average              : (*)

(*) - Time is not in sync between Sender and Reflector

#   show twamp-statistics interfaces
Interface  Last Advertisement  Delay(us)   Min(us)   Max(us)   Var(us)       Loss(%)
     xe2 2019-03-07 10:45:54       628       558       697        17   Not Enabled
```

Validation On Reflector side

```
#show running-config twamp
twamp-light reflector
 reflector-admin-state enable
 reflector-name r1 reflector-ip ipv4 30.30.30.2 reflector-port 1234 sender-ip ipv4 20.20.20.1
```

TWAMP for Client and Server with SNMP

Overview

The Two-Way Active Measurement Protocol (TWAMP) provides a standardized method for actively measuring IP network performance. It designed to evaluate metrics such as latency, jitter, and packet loss in a controlled, repeatable manner. By leveraging a client-server model, TWAMP delivers bidirectional insights that help validate service levels and troubleshoot issues proactively.

TWAMP sessions are initiated by a client, which sends control and test packets to a server (also called a reflector). The server, in turn, responds to these probes, allowing the client to calculate round-trip performance statistics.

TWAMP Client

The TWAMP client acts as the test originator. It's responsible for setting up control connections, launching measurement sessions, and collecting performance data. Each session defines sender and receiver addresses, ports, and various probe parameters, such as packet count and interval timing.

Client Feature Characteristics

Feature characteristics of the client include:

- Initiating multiple test sessions across different server addresses.
- Supporting both IPv4 and IPv6 traffic.
- Defining control parameters such as sender/receiver ports and DSCP markings.
- Starting and stopping sessions based on test scenarios.

TWAMP Server

The TWAMP server in OcNOS operates as a reflector. It listens for incoming control connections from external TWAMP clients and responds to test packets once a session is active. The server tracks connection and session details throughout the lifecycle of each test.

You can configure the server to enforce limits on active connections, sessions per connection, and per-client usage. It also supports session inactivity timeouts. By using the listen-vrf option, you can enable the server to operate across multiple VRFs, allowing performance measurements in segmented routing environments.

OcNOS version 6.1.0adds SNMP integration, allowing external systems to monitor TWAMP server activity in real time.

Server Feature Characteristics

Feature characteristics of the server include:

- Accepts TWAMP control connections and test sessions without initiating communication.
- Supports multiple VRFs using listen-vrf, allowing measurement across isolated routing domains.
- Configurable Resource Limits:
  - Maximum number of concurrent connections

© 2026 IP Infusion Inc. Proprietary

- Maximum sessions per connection
- Per-client session and connection limits
- Session inactivity timeouts

- Allows configuration of the listening port (default 862, configurable within the TCP dynamic port range).

SNMP for TWAMP

OcNOS version 6.1.0introduces the SNMP feature provides real-time visibility into TWAMP activity through trap notifications and MIB access. You can monitor events such as session starts, client timeouts, and threshold violations using standard SNMP tools.

SNMP Feature Characteristics

Feature characteristics of the SNMP include:

- **`snmpget`** and **`snmpwalk`** queries for TWAMP server data.
- SNMPv2-style traps for key operational events.
- NETCONF-based notifications for dynamic state changes.

Examples of trap-triggering events:

- A client establishing or disconnecting from a session.
- A session reaching its configured timeout.
- The server hitting its maximum connection or session limits.

All SNMP events follow a structured OID scheme under the IP Infusion enterprise MIB tree (.1.3.6.1.4.1.36673), ensuring compatibility with centralized network monitoring systems.

Benefits

The TWAMP Client Implementation has following benefits to network management and optimization:

- TWAMP clients adhere to standardized protocols, ensuring compatibility and interoperability with various TWAMP servers.
- SNMP traps provide immediate visibility into session events and performance issues.
- Automatic trap generation for connection limits, timeouts, and session states.
- Provides flexible deployment across virtual routing domains.

Prerequisites

Ensure that the TWAMP client has a stable and reliable network connection to the TWAMP server.

Topology

The network topology includes the interactions between TWAMP Server and TWAMP Client entities, highlighting the flow of control messages and test packets. The TWAMP Client initiates active measurement sessions by sending control messages to the TWAMP Server over the TWAMP Control Connection. The TWAMP Server, in response, configures the session and directs test packets to the TWAMP Session-Reflector, which ensures proper routing and measurement. The results of these measurements are then relayed back to the TWAMP Client for analysis.

© 2026 IP Infusion Inc. Proprietary

The following diagram example shows 2 nodes connected by a physical link through interfaces xe14 on each end. We will configure the first node as a TWAMP client. The second node must have a TWAMP server running.

Figure 192. TWAMP Client-Server

In the below example configurations, configure the Client for IPv4 and then for IPv6.

Server Configuration

Follow the steps below to configure the TWAMP server. These steps allow you to set up server parameters such as port number, session limits, timeouts, and VRF support.

1. Enter TWAMP server configuration mode.
```
server(config)# twamp server
```
2. Configure the listening port to define TCP port on which the TWAMP server listens for incoming connections.
```
server(config-twamp-server)# port 5555
```
3. Set Maximum Connection Duration to specify the maximum duration (in hours) that a connection remains active.
```
server(config-twamp-server)# max-connection-duration 3
```
4. Define maximum number of connections to set the total number of connections that the TWAMP server can accept concurrently.
```
server(config-twamp-server)# max-connections 17
```
5. Set maximum connections per client to limit the number of concurrent connections that a single client can establish.
```
server(config-twamp-server)# max-connections-per-client 3
```
6. Define maximum sessions to specify the total number of TWAMP sessions the server can support across all connections.
```
server(config-twamp-server)# maximum-sessions 25
```
7. Set maximum sessions per connection to limit the number of TWAMP sessions allowed per connection.
```
server(config-twamp-server)# maximum-sessions-per-connection 20
```
8. Configure inactivity timeout to specify the timeout period (in minutes) after which an inactive session will be terminated.
```
server(config-twamp-server)# inactivity-timeout 5
```
9. Add listening VRF for the active twamp server.
```
 server(config-twamp-server)# listen-vrf vrf123
```

Note: If no VRF is specified, the server runs in the default VRF.

10Enable the TWAMP Server with the configured parameters.

```
server(config-twamp-server)# server-admin-state enable
```

Validation

Show twamp server sessions information

```
ocnos-twamp#show twamp server session
Connection IDSession Id
SenderReceiverState
199 2B.00.00.03:E7548F41C6BC5154:5D180E4F
43.0.0.243.0.0.3Started
200 29.00.00.03:E7548F41C6DAA92E:103A923B
41.0.0.241.0.0.3Started
201 2C.00.00.03:E7548F41C6E5ABC0:36294112
44.0.0.244.0.0.3Started
202 2C.00.00.03:E7548F41C7161A1D:63CFE863
44.0.0.244.0.0.3Started
203 2A.00.00.03:E7548F41C82B94D9:47C8C30A
42.0.0.242.0.0.3Started
204 2C.00.00.03:E7548F41C917A461:E0C3DE71
44.0.0.244.0.0.3Started
7047-TWAMP#
```

```
OcNOS# show twamp server session Connection ID Session Id Sender Receiver State
1 5B.5B.5B.03:E00FD9FF99B7F173:D8CE5E5D 91.91.91.2
 91.91.91.3 Started
2 5B.5B.5B.03:E00FCF8221DF548E:EF824500 9.1.9.1
 91.91.91.3 Started
```

```
OcNOS# show twamp server session 5B.5B.5B.03:E00FCF8221DF548E:EF824500 State : Started
Connection ID : 2
Session ID : 5B.5B.5B.03:E00FCF8221DF548E:EF824500
Sender address : 9.1.9.1 Sender port : 44566 Receiver address: 91.91.91.3 Receiver port : 44001
Receiver VRF : default
Start time : 2022-08-31 16:44:00Timeout : 5 sec
```

Configuration to Trigger TWAMP Traps

Prerequisites

For configurations using the server and client, the following settings are enabled:

- **hardware-profile filter twamp-ipv4 enable**
- **IPv6**
- **snmp-server enable traps twamp**

For configurations using the controller and reflector, the following settings are disabled:

- **hardware-profile filter twamp-ipv4 enable and IPv6**
- **snmp-server enable traps twamp**

Basic SNMP config are already present.

- OID for IGP-TE: 1.3.6.1.4.1.36673.1600
- OID for TWAMP: 1.3.6.1.4.1.36673.1601
- OID for TWAMP-CLIENT: 1.3.6.1.4.1.36673.1603

© 2026 IP Infusion Inc. Proprietary

Perform the following steps on the TWAMP server and client devices.

1. Configure the TWAMP Server

```
server(config)# twamp-light reflector
server(config-twamp-reflector)# reflector-admin-state enable
server(config-twamp-reflector)# reflector-name R1
server(config-twamp-reflector)# reflector-ip ipv4 192.168.2.1
server(config-twamp-reflector)# reflector-port 862
server(config-twamp-reflector)# commit
server(config-twamp-reflector)# exit
server(config)# twamp server
server(config-twamp-server)# port 862
server(config-twamp-server)# max-connections 2
server(config-twamp-server)# max-connections-per-client 2
server(config-twamp-server)# maximum-sessions 64
server(config-twamp-server)# maximum-sessions-per-connection 64
server(config-twamp-server)# inactivity-timeout 1
server(config-twamp-server)# listen-vrf default
server(config-twamp-server)# server-admin-state enable
server(config-twamp-server)# commit
server(config-twamp-server)# exit
```

2. Configure the TWAMP Client.

```
client# configure terminal
client(config)# twamp-light control
client(config-twamp-control)# control-admin-state enable
client(config-twamp-control)# exit
client(config)# twamp client
client(config-twamp-client)# maximum-connections 10
client(config-twamp-client)# maximum-sessions 10
client(config-twamp-client)# connection test1
client(config-twamp-connection)# server ipv4 192.168.2.1 port 862
client(config-twamp-connection)# session s1 sender-port 1025 receiver-   port 1025 \sender-
address ipv4 192.168.2.2 receiver-address ipv4 192.168.2.1
client(config-twamp-connection)# exit-connection
client(config-twamp-client)# commit
client(config-twamp-client)# exit
```

3. Perform reflector and controller configurations for delay and loss measurement:

   a. Enable TWAMP Light Control

```
client(config)# twamp-light control
client(config-twamp-light-ctrl)# control-admin-state enable
client(config-twamp-light-ctrl)# commit
client(config-twamp-light-ctrl)# end
```

   b. Configure Delay Profile Parameters

```
client(config)# delay-profile interfaces
client(config-dp-intf)# burst-interval 2000
client(config-dp-intf)# burst-count 5
client(config-dp-intf)# interval 60
client(config-dp-intf)# advertisement accelerated
client(config-dp-intf)# advertisement accelerated threshold 10
client(config-dp-intf)# advertisement accelerated minimum-change 1
client(config-dp-intf)# notification enabled
client(config-dp-intf)# notification threshold 1
client(config-dp-intf)# notification minimum-change 10
```

   c. Configure Interface for Delay and Loss Measurement

```
client(config)# interface xe16
client(config-if)# delay-measurement dynamic twamp reflecor-ip      10.1.1.2 sender-ip
10.1.1.1
client(config-if)# loss-measurement dynamic
client(config-if)# commit
```

© 2026 IP Infusion Inc. Proprietary

1. Reflector Configuration:

   a. Enable TWAMP light reflector

```
server(config)# twamp-light reflector
server(config-twamp-light-ref)# reflector-admin-state enable
```

   b. Configure Reflector Parameters:

```
server(config-twamp-light-ref)# reflector-name r1 reflector-ip ipv4      10.1.1.2 reflector-
port 862 sender-ip ipv4 10.1.1.1 sender-port 862
server(config-twamp-light-ref)# commit
server(config-twamp-light-ref)# end
```

2. Trigger the Desired Trap.

Trap Name/OID	Twamp Server Config	Twamp Client Config	Trigger
twamp-server-client-connected or iso.3.6.1.4.1.36673.1601.0.1	CONFIG FROM CLIENT AND SERVER	Config from client and server	When a TWAMP Client connects to the TWAMP Server, the trap is generated by the TWAMP Server.
twamp-server-client-disconnected or iso.3.6.1.4.1.36673.1601.0.2	Config from client and server	Config from client and server	When the TWAMP Client removes a connection. For example: CLIENT(config-twamp-client)# no connection TEST1 The trap is generated by the TWAMP Server.
twamp-server-client-accepted iso.3.6.1.4.1.36673.1601.0.3	Config from client and server	Config from client and server	When the TWAMP Client is accepted by the TWAMP Server. The trap is generated by the TWAMP Server
twamp-server-sessions-started or iso.3.6.1.4.1.36673.1601.0.4	Config from client and server	Config from client and server	When the TWAMP Client starts an TWAMP session that is in Accepted state. Example: CLIENT# twamp client start connection CONN The trap is generated by TWAMP Server.
twamp-server-sessions-stopped or iso.3.6.1.4.1.36673.1601.0.5	Config from client and server	Config from client and server	When the TWAMP Client stops an TWAMP session that is already running.

© 2026 IP Infusion Inc. Proprietary

Trap Name/OID	Twamp Server Config	Twamp Client Config	Trigger
			Example: CLIENT# twamp client stop connection CONN The trap is generated by TWAMP Server.
twamp-server-client-inactivity-timeout or iso.3.6.1.4.1.36673.1601.0.6	inactivity-timeout 1	Config from client and server	When the TWAMP Client is connected to the TWAMP Server, but TWAMP Client did not do anything else, such as starting a TWAMP session. After the 1 minute of inactivity, the trap is generated by the TWAMP Server. And the TWAMP Client is disconnected.
twamp-server-client-maximum-connection-duration-timeout iso.3.6.1.4.1.36673.1601.0.7	maximum-connection	Already config for two connection to server exists	When the TWAMP Client reached the maximum connection duration with the TWAMP Server. The trap is generated by the TWAMP Server.
twamp-server-maximum-connections-exceeded or iso.3.6.1.4.1.36673.1601.0.8	max-connections 2	Already config for two connection to server exists	When the TWAMP Client creates the third connection, the TWAMP Server will not accept it, and the TWAMP server will generate the trap.
twamp-server-maximum-sessions-exceeded or iso.3.6.1.4.1.36673.1601.0.9	maximum-sessions 2	Already config for two session to server exists	When the TWAMP Client creates the third session the TWAMP Server will not accept it, and the TWAMP server will generate the trap.
twamp-server-maximum-connections-per-client-exceeded or iso.3.6.1.4.1.36673.1601.0.10	max-connections-per-client 2	Already two connection to sever exists	When the TWAMP Client creates the third connection, the TWAMP Server will not accept it, and the TWAMP server will generate the trap.
twamp-server-maximum-sessions-	maximum-sessions-	Already 1 session to	When the TWAMP Client

Trap Name/OID	Twamp Server Config	Twamp Client Config	Trigger
per-connection-exceeded or iso.3.6.1.4.1.36673.1601.0.11	per-connection 1	server exists	creates the third connection, the TWAMP Server will not accept it, and the TWAMP server will generate the trap.
twamp-delay-threshold-breached or.1.3.6.1.4.1.36673.1601.0.12	Config from controller and reflector	Config from controller and reflector	When the TWAMP Client creates the second session the TWAMP Server will not accept it, and the TWAMP server will generate the trap. These traps are generated by the TWAMP Client.
twamp-client-session-accepted or 1.3.6.1.4.1.36673.1603.0.1	Config from client and server	Config from client and server	When a Session is accepted by the TWAMP Server to the TWAMP Client. So, The TWAMP Client can get to know when the acceptance as been done, since the accept may not be synchronous. These traps are generated by the TWAMP Client.
twamp-client-sessions-started or .1.3.6.1.4.1.36673.1603.0.2	Config from client and server	Config from client and server	When a Session is started with the Twamp Server, the trap is generated to inform the TWAMP Client side that it has started. These traps are generated by the TWAMP Client.
twamp-client-sessions-stopped or .1.3.6.1.4.1.36673.1603.0.3	Config from client and server	Config from client and server	When a Session is stopped with the Twamp Server, the trap is

© 2026 IP Infusion Inc. Proprietary

Trap Name/OID	Twamp Server Config	Twamp Client Config	Trigger
			generated to inform the TWAMP Client side that it has stopped These traps are generated by the TWAMP Client.
twamp-delay-measurement-threshold-crossed or .1.3.6.1.4.1.36673.1600.0.1	Config from controller and reflector	Config from controller and reflector	When an interface is configured >to run the interface delay profile, it will calculate the RTT delay. If the delay suppress the default thresholds (in case no thresholds were set) or if it crosses the configured thresholds, the trap will be generated. These traps are generated by the TWAMP Client.
twamp-delay-measurement-loss-changed or .1.3.6.1.4.1.36673.1600.0.2	Config from controller and reflector D	Config from controller and reflector	When the packet loss changes from the last measured loss this trap will be generated. These traps are generated by the TWAMP Client.
twamp-overall-rate-crossed or .1.3.6.1.4.1.36673.1601.0.13	Use scale config for exceeding the hardware limit (reflector+sender) for Q1-64 Q2-512	Use scale config for exceeding the hardware limit (reflector+sender) for Q1-64 Q2-512	When the overall supported TWAMP rate limit has been crossed.
twamp-overall-rate-crossed-clear or .1.3.6.1.4.1.36673.1601.0.14	Reduce the number of session below threshold after the scale config are attained for	Reduce the number of session below threshold after the scale config are attained for	When the overall crossed TWAMP rate limit has been cleared.

Trap Name/OID	Twamp Server Config	Twamp Client Config	Trigger
	Q1-64 Q2-512	Q1-64 Q2-512	

Validation

TRAPS: (Highlighted)

For **twamp-server-client-connected and twamp-server-session-accepted**

```
iso.3.6.1.2.1.1.3.0 = Timeticks: (3523800) 9:47:18.00 iso.3.6.1.6.3.1.1.4.1.0 = OID:
iso.3.6.1.4.1.36673.1601.0.1 iso.3.6.1.2.1.1.5.0 = STRING: "TWAMPSERVER"
iso.3.6.1.4.1.36673.1601.9.1.1 = Gauge32: 4 iso.3.6.1.4.1.36673.1601.9.1.2 = STRING: "192.168.2.2"
iso.3.6.1.4.1.36673.1601.9.1.3 = Gauge32: 21191 iso.3.6.1.4.1.36673.1601.9.1.4 = STRING: "default"
2025-04-12 20:29:52 <UNKNOWN> [UDP: [10.16.113.127]:46007->[10.16.99.107]:12162]:
iso.3.6.1.2.1.1.3.0 = Timeticks: (3523800) 9:47:18.00 iso.3.6.1.6.3.1.1.4.1.0 = OID:
iso.3.6.1.4.1.36673.1601.0.3
NOTIFICATION:
Incoming notification:
<notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
<eventTime>2025-04-12T16:16:02Z</eventTime>
<twamp-server-client-connected xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-twamp">
<severity>minor</severity>
<eventClass>state</eventClass>
<connection-id>5</connection-id>
<client-address>192.168.2.2</client-address>
<client-port>33482</client-port>
<client-vrf>default</client-vrf>
</twamp-server-client-connected>
</notification>
Incoming notification:
<notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
<eventTime>2025-04-12T16:16:02Z</eventTime>
<twamp-server-session-accepted xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-twamp">
<severity>minor</severity>
<eventClass>state</eventClass>
<session-id>C0.A8.02.01:EBA510428393CCD0:B92D265A</session-id>
<connection-id>6</connection-id>
<sender-address>192.168.2.2</sender-address>
<sender-port>1025</sender-port>
<receiver-address>192.168.2.1</receiver-address>
<receiver-port>1025</receiver-port>
<receiver-vrf>default</receiver-vrf>
</twamp-server-session-accepted>
</notification>
```

twamp-server-client-”disconnected” or “iso.3.6.1.4.1.36673.1601.0.2”

```
iso.3.6.1.2.1.1.3.0 = Timeticks: (4422700) 12:17:07.00 iso.3.6.1.6.3.1.1.4.1.0 = OID:
iso.3.6.1.4.1.36673.1601.0.2 iso.3.6.1.2.1.1.5.0 = STRING: "TWAMPSERVER"
iso.3.6.1.4.1.36673.1601.9.1.1 = Gauge32: 9 iso.3.6.1.4.1.36673.1601.9.1.2 = STRING: "192.168.2.2"
iso.3.6.1.4.1.36673.1601.9.1.3 = Gauge32: 37 iso.3.6.1.4.1.36673.1601.9.1.4 = STRING: "default"
NOTIFICATIONS:
Incoming notification:
<notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
<eventTime>2025-04-12T17:33:00Z</eventTime>
<twamp-server-client-disconnected xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-twamp">
<severity>minor</severity>
<eventClass>state</eventClass>
<connection-id>11</connection-id>
<client-address>192.168.2.2</client-address>
<client-port>37</client-port>
<client-vrf>default</client-vrf>
</twamp-server-client-disconnected>
</notification>
```

© 2026 IP Infusion Inc. Proprietary

twamp-server-sessions-started" or "iso.3.6.1.4.1.36673.1601.0.4

Start connection from client to generate traps

TRAPS:

```
iso.3.6.1.2.1.1.3.0 = Timeticks: (4483300) 12:27:13.00 iso.3.6.1.6.3.1.1.4.1.0 = OID:
iso.3.6.1.4.1.36673.1601.0.4 iso.3.6.1.2.1.1.5.0 = STRING:
"TWAMPSERVER" iso.3.6.1.4.1.36673.1601.10.1.1 = STRING: "C0.A8.02.01:EBA522DDC5B9D7FD:8C0A1D0A"
iso.3.6.1.4.1.36673.1601.10.1.2 = Gauge32: 12
NOTIFICATION:
Incoming notification:
<notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
<eventTime>2025-04-12T17:44:37Z</eventTime>
<twamp-server-sessions-started xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-twamp">
<severity>minor</severity>
<eventClass>state</eventClass>
<session-id>C0.A8.02.01:EBA524F84EA92E62:D5E78A33</session-id>
<connection-id>13</connection-id>
</twamp-server-sessions-started>
</notification>
```

twamp-server-sessions-stopped" or "iso.3.6.1.4.1.36673.1601.0.5

TRAPS:

```
iso.3.6.1.2.1.1.3.0 = Timeticks: (4645600) 12:54:16.00 iso.3.6.1.6.3.1.1.4.1.0 = OID:
iso.3.6.1.4.1.36673.1601.0.5 iso.3.6.1.2.1.1.5.0 = STRING: "TWAMPSERVER"
iso.3.6.1.4.1.36673.1601.10.1.1 = STRING: "C0.A8.02.01:EBA52644088626F6:FD04027E"
iso.3.6.1.4.1.36673.1601.10.1.2 = Gauge32: 14
2025-04-12 23:36:49 <UNKNOWN> [UDP: [10.16.113.121]:39643->[10.16.99.107]:12162]:
NOTIFICATION:
Incoming notification:
<notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
<eventTime>2025-04-12T18:06:41Z</eventTime>
<twamp-server-sessions-stopped xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-twamp">
<severity>minor</severity>
<eventClass>state</eventClass>
<session-id>C0.A8.02.01:EBA52644088626F6:FD04027E</session-id>
<connection-id>14</connection-id>
</twamp-server-sessions-stopped>
</notification>
```

twamp-server-client-inactivity-timeout or iso.3.6.1.4.1.36673.1601.0.6

TRAPS: (Highlighted)

```
iso.3.6.1.2.1.1.3.0 = Timeticks: (4763300) 13:13:53.00 iso.3.6.1.6.3.1.1.4.1.0 = OID:
iso.3.6.1.4.1.36673.1601.0.6 iso.3.6.1.2.1.1.5.0 = STRING: "TWAMPSERVER"
iso.3.6.1.4.1.36673.1601.9.1.1 = Gauge32: 15 iso.3.6.1.4.1.36673.1601.9.1.7 = Hex-STRING: 07 E9 04 0C
12 18 1D 00
NOTIFICATION:
yangcli ocnos@10.16.113.127> 2025 Apr 12 18:32:18.887 : TWAMPSERVER : OAM : WARN : Connection 17 has
been inactive for 1 minutes, closing the conenction
Incoming notification:
<notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
<eventTime>2025-04-12T18:40:18Z</eventTime>
<twamp-server-client-inactivity-timeout xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-twamp">
<severity>minor</severity>
<eventClass>state</eventClass>
<connection-id>21</connection-id>
<last-activity-time>2025-04-12T18:38:23Z</last-activity-time>
</twamp-server-client-inactivity-timeout>
</notification>
```

twamp-server-maximum-connections-exceeded or iso.3.6.1.4.1.36673.1601.0.8

TRAPS: (Highlighted)

```
2025-04-13 00:51:22 <UNKNOWN> [UDP: [10.16.113.127]:46007->[10.16.99.107]:12162]:
iso.3.6.1.2.1.1.3.0 = Timeticks: (5092800) 14:08:48.00 iso.3.6.1.6.3.1.1.4.1.0 = OID:
iso.3.6.1.4.1.36673.1601.0.8 iso.3.6.1.2.1.1.5.0 = STRING: "TWAMPSERVER" iso.3.6.1.4.1.36673.1601.1 =
Gauge32: 2
```

© 2026 IP Infusion Inc. Proprietary

```
NOTIFICATION:
yangcli ocnos@10.16.113.127> 2025 Apr 12 19:27:33.057 : TWAMPSERVER : OAM : ERROR : Maximum number of
connections 2 reached, rejecting the new connection
Incoming notification:
<notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
<eventTime>2025-04-12T19:27:33Z</eventTime>
<twamp-server-maximum-connections-exceeded xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-twamp">
<severity>minor</severity>
<eventClass>state</eventClass>
<max-connections>2</max-connections>
</twamp-server-maximum-connections-exceeded>
</notification>
```

twamp-server-maximum-sessions-exceeded" or "iso.3.6.1.4.1.36673.1601.0.9"
TWAMPSERVER CONFIG:

```
TWAMPSERVER#sh run twamp
twamp-light reflector
reflector-admin-state enable
reflector-name R1 reflector-ip ipv4 192.168.2.1 reflector-port 862
twamp server
port 862
max-connections 32
max-connections-per-client 4
maximum-sessions 2
maximum-sessions-per-connection 1
listen-vrf default
server-admin-state enable
TRAPS: (Highlighted)
iso.3.6.1.2.1.1.3.0 = Timeticks: (15810800) 1 day, 19:55:08.00 iso.3.6.1.6.3.1.1.4.1.0 = OID:
iso.3.6.1.4.1.36673.1601.0.9
iso.3.6.1.2.1.1.5.0 = STRING: "TWAMPSERVER" iso.3.6.1.4.1.36673.1601.3 = Gauge32: 2
NOTIFICATION:
yangcli ocnos@10.16.113.127> create-subscription
RPC OK Reply 1 for session 11:
yangcli ocnos@10.16.113.127> 2025 Apr 14 01:13:31.831 : TWAMPSERVER : OAM : WARN : Maximum number of
sessions 2 reached on the server
Incoming notification:
<notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
<eventTime>2025-04-14T01:13:31Z</eventTime>
<twamp-server-client-connected xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-twamp">
<severity>minor</severity>
<eventClass>state</eventClass>
<connection-id>54</connection-id>
<client-address>192.168.2.2</client-address>
<client-port>44259</client-port>
<client-vrf>default</client-vrf>
</twamp-server-client-connected>
</notification>
Incoming notification:
<notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
<eventTime>2025-04-14T01:13:31Z</eventTime>
<twamp-server-maximum-sessions-exceeded xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-twamp">
<severity>minor</severity>
<eventClass>state</eventClass>
<maximum-sessions>2</maximum-sessions>
</twamp-server-maximum-sessions-exceeded>
</notification>
```

twamp-server-maximum-connections-per-client-exceeded" or "iso.3.6.1.4.1.36673.1601.0.10"
TRAPS: (Highlighted)

```
iso.3.6.1.2.1.1.3.0 = Timeticks: (16085400) 1 day, 20:40:54.00 iso.3.6.1.6.3.1.1.4.1.0 = OID:
iso.3.6.1.4.1.36673.1601.0.10 iso.3.6.1.2.1.1.5.0 = STRING: "TWAMPSERVER" iso.3.6.1.4.1.36673.1601.2
= Gauge32: 2 iso.3.6.1.4.1.36673.1601.9.1.1 = Gauge32: 59
NOTIFICATIONS:
yangcli ocnos@10.16.113.127> 2025 Apr 14 01:57:50.958 : TWAMPSERVER : OAM : ERROR : Maximum number
connections 2 from client 192.168.2.2 reached
Incoming notification:
```

```
<notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
<eventTime>2025-04-14T01:57:50Z</eventTime> Two-way Active Measurement Protocol Server <twamp-server-
maximum-connections-per-client-exceeded xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-twamp">
<severity>minor</severity>
<eventClass>state</eventClass>
<max-connections-per-client>2</max-connections-per-client>
<connection-id>59</connection-id>
</twamp-server-maximum-connections-per-client-exceeded>
</notification>
```

twamp-server-maximum-sessions-per-connection-exceeded" or "iso.3.6.1.4.1.36673.1601.0.11"

TRAPS: (Highlighted)

```
iso.3.6.1.2.1.1.3.0 = Timeticks: (15881200) 1 day, 20:06:52.00 iso.3.6.1.6.3.1.1.4.1.0 = OID:
iso.3.6.1.4.1.36673.1601.0.11 iso.3.6.1.2.1.1.5.0 = STRING: "TWAMPSERVER" iso.3.6.1.4.1.36673.1601.4
= Gauge32: 1 iso.3.6.1.4.1.36673.1601.9.1.1 = Gauge32: 51
NOTIFICATIONS: Two-way Active Measurement Protocol Server
yangcli ocnos@10.16.113.127> create-subscription
RPC OK Reply 1 for session 12:
yangcli ocnos@10.16.113.127> 2025 Apr 14 01:24:46.865 : TWAMPSERVER : OAM : WARN : Maximum number of
sessions 1 reached for client 51
Incoming notification:
<notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
<eventTime>2025-04-14T01:24:46Z</eventTime>
<twamp-server-maximum-sessions-per-connection-exceeded
xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-twamp">
<severity>minor</severity>
<eventClass>state</eventClass>
<maximum-sessions-per-connection>1</maximum-sessions-per-connection>
<connection-id>51</connection-id>
</twamp-server-maximum-sessions-per-connection-exceeded>
</notification>
yangcli ocnos@10.16.113.127>
```

twamp-overall-rate-crossed" or ".1.3.6.1.4.1.36673.1601.0.13

TRAPS: (Highlighted)

```
.1.3.6.1.2.1.1.3.0 = Timeticks: (351700) 0:58:37.00     .1.3.6.1.6.3.1.1.4.1.0 = OID:
.1.3.6.1.4.1.36673.1601.0.13      .1.3.6.1.2.1.1.5.0 = STRING: 7035       .1.3.6.1.4.1.36673.1601.8 =
Gauge32: 335 pps     .1.3.6.1.4.1.36673.1601.9 = Gauge32: 310 pps    .1.3.6.1.4.1.36673.1601.10 =
Gauge32: 25 pps    .1.3.6.1.4.1.36673.1601.11 = Gauge32: 320 pps   .1.3.6.1.4.1.36673.1601.12 =
INTEGER: true(1)

Incoming notification:
notification {
  eventTime 2026-02-05T15:56:41.831Z
  twamp-overall-rate-crossed {
    severity major
    eventClass state
    packet-sending-rate 340
    senders-rate 315
    reflectors-rate 25
    max-sending-rate 320
    max-sending-rate-crossed true
  }
}
```

twamp-overall-rate-crossed-clear" or ".1.3.6.1.4.1.36673.1601.0.14

TRAPS: (Highlighted)

```
.1.3.6.1.2.1.1.3.0 = Timeticks: (365800) 1:00:58.00     .1.3.6.1.6.3.1.1.4.1.0 = OID:
.1.3.6.1.4.1.36673.1601.0.14      .1.3.6.1.2.1.1.5.0 = STRING: 7035       .1.3.6.1.4.1.36673.1601.8 =
Gauge32: 315 pps     .1.3.6.1.4.1.36673.1601.9 = Gauge32: 310 pps    .1.3.6.1.4.1.36673.1601.10 =
Gauge32: 5 pps     .1.3.6.1.4.1.36673.1601.11 = Gauge32: 320 pps   .1.3.6.1.4.1.36673.1601.12 =
INTEGER: false(0)

Incoming notification:
```

© 2026 IP Infusion Inc. Proprietary

```
notification {
  eventTime 2026-02-05T15:57:11.876Z
  twamp-overall-rate-crossed-cleared {
    severity warning
    eventClass state
    packet-sending-rate 315
    senders-rate 315
    reflectors-rate 0
    max-sending-rate 320
    max-sending-rate-crossed false
  }
}
```

Glossary

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

TWAMP	Two-Way Active Measurement Protocol, a standardized protocol for measuring network performance in a bidirectional manner.
Client	In the context of TWAMP, the client is a device or software component responsible for initiating TWAMP test sessions and measuring network performance.
Session	A measurement session in TWAMP involves the exchange of test packets between the client and the server to assess network performance.
Server	The counterpart to the client, the server in TWAMP receives test packets, conducts measurements, and provides results to the client.
Test Packets	Data packets exchanged between the client and server to measure network performance, including metrics like latency, jitter, and packet loss.
Initiate	The action of starting a TWAMP test session from the client to the server.
Latency	The time it takes for data packets to travel from the client to the server and back, often measured in milliseconds.

© 2026 IP Infusion Inc. Proprietary

TWAMP for VPN Inter-Access Interface

Overview

The Two-Way Active Measurement Protocol (TWAMP), defined by RFC 5357, measures end-to-end network performance metrics, such as one-way or two-way latency (delay) and packet loss, between two network endpoints. Service providers and enterprises commonly use TWAMP to verify Service Level Agreements (SLAs) and diagnose network performance issues.The TWAMP for VPN Inter-Access Interface feature extends TWAMP functionality to multi-VPN or multi-VRF environments. It enables end-to-end performance measurements across access interfaces of network endpoints. This capability is particularly valuable in service provider networks, where multiple customers or internal groups operate within isolated, logically separated virtual private networks (VPNs), yet require cross-boundary performance metrics.

Feature Characteristics

When configured for VPN Inter-Access Interfaces (for example, PE1 interface xe4 and PE2 interface xe14), TWAMP test packets traverse the specified path, adhering to configured Quality of Service (QoS) policies, queuing mechanisms, and other interface settings. This ensures that measurements accurately reflect the actual user-experienced delay between access interfaces.

Benefits

TWAMP for VPN Inter-Access Interface has following benefits to network management and optimization:

- Enables delay measurements across VPN Inter-Access Interfaces to verify network SLAs.
- Supports troubleshooting of performance issues in multi-VPN environments.

Prerequisites

- To enable TWAMP measurements on VPN Inter-Access Interfaces, configure the following hardware-profile filters:

```
OcNOS(config)# hardware-profile filter egress-ipv4 enable
OcNOS(config)# hardware-profile filter twamp-ipv4 enable
```

Note: A system reboot is required to apply these hardware filters.

Limitations

- The inter-access VPN TWAMP session does not support multi-path scenarios when traffic flows through an interface other than the vpn-net-interface. If TWAMP packets are received on a different interface, the system treats them as invalid, which increases the packet loss rate.
- Inter-access VPN TWAMP is a Layer 3 feature and is not designed to operate on LAG member interfaces.

© 2026 IP Infusion Inc. Proprietary

Topology

The topology includes two Customer Edge (CE) devices, each connected to a Provider Edge (PE) device (PE1 and PE2). The PEs are interconnected through an MPLS core that supports Layer 3 VPN (L3VPN) services with a VPN tunnel. For simplicity, the VPN cloud is not shown, and PE1 and PE2 are depicted as directly connected.

Figure 193. TWAMP for VPN inter-access interface

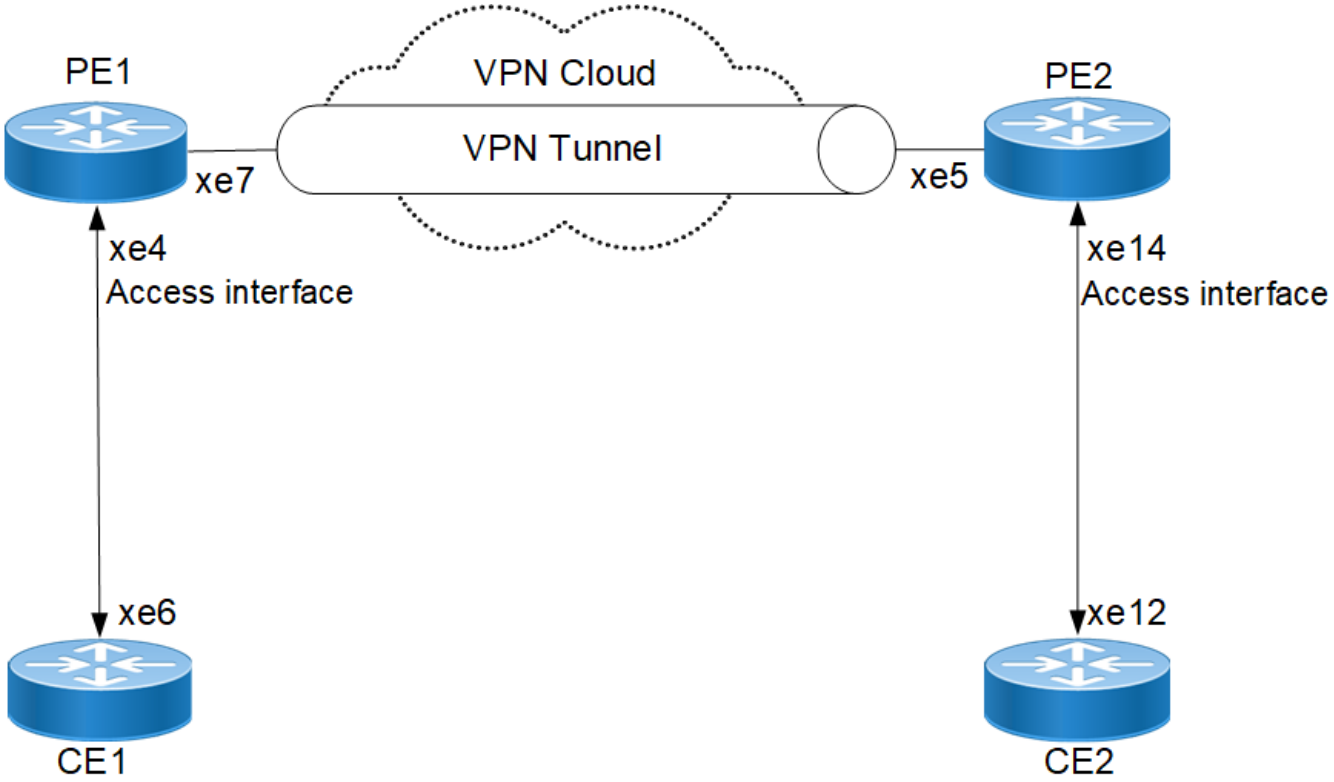

Configurations

PE1

1. Enter Configuration Mode

   a. Access the configuration terminal to begin modifying the router settings.

   ```
   OcNOS#configure terminal
   ```

   b. Configure VRF for VPN (vrf101):
   Set up a Virtual Routing and Forwarding (VRF) instance named vrf101, assign a Route Distinguisher (RD), and Route Target (RT).

   ```
   (config)#ip vrf vrf101
   (config-vrf)#rd 1.1.1.1:101
   (config-vrf)#route-target both 65010:101
   (config-vrf)#exit
   (config)#commit
   ```

   c. Configure LDP
   Enter Label Distribution Protocol (LDP) mode and set the transport address to the loopback IP.

© 2026 IP Infusion Inc. Proprietary

```
(config)#router ldp
(config-router-ldp)#transport-address ipv4 1.1.1.1
(config-router-ldp)#exit
```

d. Configure Loop-back Interfaces
Set up loop-back interfaces for global and VRF contexts with primary and secondary IP addresses.

```
(config)#interface lo
(config-if)#ip address 1.1.1.1/32 secondary
(config-if)#exit
```

e. Configure Access Interface (xe4)
Bind the access interface to vrf101 and assign an IP address for connectivity to the CE device.

```
(config)#interface xe4
(config-if)#ip vrf forwarding vrf101
(config-if)#ip address 103.1.1.1/24
(config-if)#description ACCESS-INTERFACE
(config-if)#exit
```

f. Configure Transport Interface (xe7)
Configure the transport interface with an IP address, enable label switching, and activate LDP for MPLS.

```
(config)# interface xe7
(config-if)# ip address 11.1.1.1/24
(config-if)# label-switching
(config-if)# enable-ldp ipv4
(config-if)# description TRANSPORT-INTERFACE
(config-if)# exit
(config)# commit
```

g. Configure OSPF
Set up OSPF for underlay routing, specifying the router ID and networks to participate in OSPF area 0.

```
(config)# router ospf 1
(config-router-ospf)# ospf router-id 1.1.1.1
(config-router-ospf)# network 1.1.1.1/32 area 0.0.0.0
(config-router-ospf)# network 11.1.1.0/24 area 0.0.0.0
(config-router-ospf)# exit
(config)# commit
```

h. Configure BGP
Set up BGP for overlay routing, including iBGP peering with PE2 and EVPN and VRF address families.

```
(config)# router bgp 65010
(config-router-bgp)# bgp router-id 1.1.1.1
(config-router-bgp)# neighbor 3.3.3.3 remote-as 65010
(config-router-bgp)# neighbor 3.3.3.3 update-source lo

(config-router-bgp)# address-family vpnv4 unicast
(config-router-bgp-af)# neighbor 3.3.3.3 activate
(config-router-bgp-af)# neighbor 3.3.3.3 next-hop-self
(config-router-bgp-af)# exit-address-family

(config-router-bgp)# address-family ipv4 vrf vrf101
(config-router-bgp-vrf)# redistribute connected
(config-router-bgp-vrf)# exit-address-family

(config-router-bgp)# exit
(config)# commit
```

i. Enable TWAMP-Light Control
Enable the TWAMP-light administrative state to allow TWAMP operations.

```
(config)# twamp-light control
(config-twamp-light-ctrl)# control-admin-state enable
(config-twamp-light-ctrl)# exit
```

© 2026 IP Infusion Inc. Proprietary

```
(config)# commit
```

j. Reload the System
Save the configuration and reload the system to apply hardware filter changes.

```
(config)# exit
OcNOS# reload
The system has unsaved changes.
Would you like to save them now? (y/n): y
Are you sure you would like to reset the system? (y/n): y
```

k. TWAMP Session
Create a TWAMP session for VPN inter-access interface measurement.

```
OcNOS# configure terminal
(config)# twamp session PE1-SESSION
(config-twamp-session)# sender-ip 103.1.1.1
(config-twamp-session)# sender-port 862
(config-twamp-session)# reflector-ip 203.1.1.1
(config-twamp-session)# reflector-port 862
(config-twamp-session)# vrf-name vrf101
(config-twamp-session)# vpn-net-interface xe7
(config-twamp-session)# description TWAMP VPN Inter-Acc-Intf
(config-twamp-session)# dscp 0
(config-twamp-session)# burst-interval 2000
(config-twamp-session)# burst-count 2
(config-twamp-session)# interval 35
(config-twamp-session)# notification enabled
(config-twamp-session)# notification threshold 80
(config-twamp-session)# notification minimum-change 80
(config-twamp-session)# no shut
(config-twamp-session)# exit
(config)# commit
```

PE2

1. Enter Configuration Mode

a. Access the configuration terminal to begin modifying the router settings.

```
OcNOS# configure terminal
```

b. Configure VRF for VPN (vrf101)
Create a VRF instance named vrf101, assign an RD, and RT.

```
(config)# ip vrf vrf101
(config-vrf)# rd 3.3.3.3:101
(config-vrf)# route-target both 65010:101
(config-vrf)# exit
(config)# commit
```

c. Configure LDP
Enter LDP mode and set the transport address to the loopback IP.

```
(config)# router ldp
(config-router-ldp)# transport-address ipv4 3.3.3.3
(config-router-ldp)# exit
```

d. Configure Loop-back Interfaces
Set up loopback interfaces for global and VRF contexts with primary and secondary IP addresses.

```
(config)# interface lo
(config-if)# ip address 3.3.3.3/32 secondary
(config-if)# exit
```

e. Configure Access Interface (xe14)
Bind the access interface to vrf101 and assign an IP address.

© 2026 IP Infusion Inc. Proprietary

```
(config)# interface xe14
(config-if)# ip vrf forwarding vrf101
(config-if)# ip address 203.1.1.1/24
(config-if)# description ACCESS-INTERFACE
(config-if)# exit
```

f. Configure Transport Interface (xe5)
Configure the transport interface with an IP address, enable label switching, and activate LDP for MPLS.

```
(config)# interface xe5
(config-if)# ip address 11.1.1.2/24
(config-if)# label-switching
(config-if)# enable-ldp ipv4
(config-if)# description NETWORK-INTERFACE
(config-if)# exit
(config)# commit
```

g. Configure OSPF
Set up OSPF for underlay routing, specifying the router ID and networks to participate in OSPF area 0.

```
(config)# router ospf 1
(config-router-ospf)# ospf router-id 3.3.3.3
(config-router-ospf)# network 3.3.3.3/32 area 0.0.0.0
(config-router-ospf)# network 11.1.1.0/24 area 0.0.0.0
(config-router-ospf)# exit
(config)# commit
```

h. Configure BGP
Set up BGP for overlay routing, including iBGP peering with PE1 and EVPN and VRF address families.

```
(config)# router bgp 65010
(config-router-bgp)# bgp router-id 3.3.3.3
(config-router-bgp)# neighbor 1.1.1.1 remote-as 65010
(config-router-bgp)# neighbor 1.1.1.1 update-source lo

(config-router-bgp)# address-family vpnv4 unicast
(config-router-bgp-af)# neighbor 1.1.1.1 activate
(config-router-bgp-af)# neighbor 1.1.1.1 next-hop-self
(config-router-bgp-af)# exit-address-family

(config-router-bgp)# address-family ipv4 vrf vrf101
(config-router-bgp-vrf)# redistribute connected
(config-router-bgp-vrf)# exit-address-family

(config-router-bgp)# exit
(config)# commit
```

i. Configure TWAMP-Light Reflector
Configure TWAMP reflectors for VPN inter-access interface measurement.

```
(config)# router bgp 65010
(config-router-bgp)# bgp router-id 3.3.3.3
(config-router-bgp)# neighbor 1.1.1.1 remote-as 65010
(config-router-bgp)# neighbor 1.1.1.1 update-source lo

(config-router-bgp)# address-family vpnv4 unicast
(config-router-bgp-af)# neighbor 1.1.1.1 activate
(config-router-bgp-af)# neighbor 1.1.1.1 next-hop-self
(config-router-bgp-af)# exit-address-family

(config-router-bgp)# address-family ipv4 vrf vrf101
(config-router-bgp-vrf)# redistribute connected
(config-router-bgp-vrf)# exit-address-family

(config-router-bgp)# exit
(config)# commit
```

j. Reload the System
Save the configuration and reload the system to apply hardware filter changes.

```
(config)# exit
OcNOS# reload
The system has unsaved changes.
Would you like to save them now? (y/n): y
Are you sure you would like to reset the system? (y/n): y
```

Validation

For validation and comparison, the TWAMP link measurement was configured on the interface xe7 on PE1 and the reflector on PE2 running at interface xe5. Thus, it is possible to compare the VPN Tunnel measurements and the Access to Access measurements as shown below:

To verify the TWAMP configuration on PE1, use the following commands to check the running configuration and session status:

- Check TWAMP Configuration

```
#show running-config twamp
twamp session PE1-SESSION
sender-ip 103.1.1.1
sender-port 862
reflector-ip 203.1.1.1
reflector-port 862
vrf-name vrf101
vpn-net-interface xe7
description TWAMP VPN Inter-Acc-Intf
dscp 0
burst-interval 2000
burst-count 2
interval 35
notification enabled
notification threshold 80
notification minimum-change 80
no shut
```

- Verify TWAMP Session Status

```
#show twamp-session all
=======================================
Session Name      : PE1-SESSION
VPN Network Interface name    : xe7
State             : Running
Sender IP         : 103.1.1.1
Sender Port       : 862
DSCP value        : default
Vrf-Name          : vrf101
Reflector IP      : 203.1.1.1
Reflector port    : 862
Execution Status  : Running
HW Status         : HW rules installed
Last Calculated stats:
Time: 2025-06-12 12:04:22
  Average delay          : 22
  Minimum delay          : 20
  Maximum delay          : 23
  Average delay variation: 0
  Minimum delay variation: 0
  Maximum delay variation: 0
  Packets sent           : 30
  Packets received       : 30
  Packets timeout        : 0
Packet Loss : 0.00
Notification stats:
  Notification : Enabled
  Threshold (%): 80
  Minimum-change (usec): 80
```

```
  Last Notified Packet Loss:
    Time: 2025-06-12 12:12:12
    Packet Loss: 30.00
  Last Notified Average Delay:
    Time: -
    Average Delay (usec):
```

- Check TWAMP Statistics

```
  #show twamp-statistics
========================================
      TWAMP Test-Session Statistics
========================================
  Test-Session Name        : PE1-SESSION
  Start Time               : 2025 Jun 12 11:43:20
  Elapsed time(milli sec)  : 1305917
  Packets Sent             : 1304
  Packets Received         : 1304
  Packet Loss(%)           : 0.00
  Round Trip Delay(usec)
      Minimum              : 20
      Maximum              : 24
      Average              : 21
  Forward Delay(usec)
      Minimum              : (*)
      Maximum              : (*)
      Average              : (*)
  Reverse Delay(usec)
      Minimum              : (*)
      Maximum              : (*)
      Average              : (*)
  Round Trip Delay Variation(usec)
      Minimum              : 15
      Maximum              : 24
      Average              : 15
  Forward Delay Variation(usec)
      Minimum              : (*)
      Maximum              : (*)
      Average              : (*)
  Reverse Delay Variation(usec)
      Minimum              : (*)
      Maximum              : (*)
      Average              : (*)
(*) - Time is not in sync between Sender and Reflector
```

- Check Specific TWAMP Session

```
  #show twamp-session PE1-SESSION
Session Name      : PE1-SESSION
VPN Network Interface name    : xe7
State             : Running
Sender IP         : 103.1.1.1
Sender Port       : 862
DSCP value        : default
Vrf-Name          : vrf101
Reflector IP      : 203.1.1.1
Reflector port    : 862
Execution Status  : Running
HW Status         : HW rules installed
Last Calculated stats:
Time: 2025-06-12 12:04:22
  Average delay          : 22
  Minimum delay          : 20
  Maximum delay          : 23
  Average delay variation: 0
  Minimum delay variation: 0
  Maximum delay variation: 0
  Packets sent           : 30
```

```
  Packets received       : 30
  Packets timeout        : 0
Packet Loss : 0.00
Notification stats:
  Notification : Enabled
  Threshold (%): 80
  Minimum-change (usec): 80
  Last Notified Packet Loss:
    Time: 2025-06-12 12:12:12
    Packet Loss: 30.00
  Last Notified Average Delay:
    Time: -
    Average Delay (usec): -

#show twamp-session
Session-Name  Net-Intf Delay(us) Min(us) Max(us) Var(us) Loss(%)   State
PE1-SESSION     xe7      22        20      23      0       0.00    Running
```

To verify the TWAMP configuration on PE2, use the following command to check the reflector configuration:

- Check TWAMP Reflector Configuration

```
#show running-config twamp
twamp-light reflector
  reflector-admin-state enable
  reflector-name R4 reflector-ip ipv4 203.1.1.1
reflector-port 862 sender-ip
ipv4 103.1.1.1 sender-port 862 vrf vrf101 vpn-net-interface
xe5
```

CLI Commands

This section describes the CLI commands used to configure TWAMP over VPN Inter-Access Interfaces.

© 2026 IP Infusion Inc. Proprietary

twamp session

Use this command to configure a TWAMP session for end-to-end performance measurement across VPN inter-access interfaces.

Command Syntax

```
twamp session WORD
no twamp session WORD
```

Parameters

WORD
Set TWAMP Session name

Default

None

Command Mode

None

Applicability

This command is introduced in OcNOS version 6.6.1.

Example

The following example shows the creation of a TWAMP Session and the change into `twamp-session` mode prompt.

```
OcNOS#configure terminal
OcNOS(config)#twamp session PE1-SESSION
OcNOS(twamp-session)#
```

© 2026 IP Infusion Inc. Proprietary

no twamp session all

Use this command to remove all configured TWAMP sessions in a single command. Requires all sessions to be in shut mode; otherwise, the command fails and displays an error.

Command Syntax

```
no twamp session all
```

Parameters

None

Default

None

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1.

Example

The following example shows the deletion of all TWAMP Sessions via the single-command running on configure-mode.

```
OcNOS#configure terminal
OcNOS(config)# no twamp session all
```

shut

Use this command to shut a configured TWAMP Session. It will stop the TWAMP Session execution and measurement.

Command Syntax

```
shut
```

Parameters

None

Default

Enabled

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1.

Example

The following example shows the TWAMP Session being set to shut, which will then stop the TWAMP Session named PE1-SESSION.

```
OcNOS#configure terminal
OcNOS(config)#twamp session PE1-SESSION
OcNOS(twamp-session)#shut
```

© 2026 IP Infusion Inc. Proprietary

no shut

Use this command to unshut a configured TWAMP Session. It will validate the configuration and start the TWAMP Session execution and measurement when the configuration is correct. Otherwise, the respective error is thrown related to the configuration validation that caused the error.

Command Syntax

```
no shut
```

Parameters

None

Default

Disable

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1.

Example

The following example shows the TWAMP Session being set to no shut, which will then validate and start the TWAMP Session named PE1-SESSION.

```
OcNOS#configure terminal
OcNOS(config)#twamp session PE1-SESSION
OcNOS(twamp-session)# no shut
```

© 2026 IP Infusion Inc. Proprietary

reflector-ip

Use this command to set the IP of the TWAMP Reflector so that this TWAMP Session will send the test packets.

Command Syntax

```
reflector-ip (A.B.C.D|HOSTNAME)
```

Parameters

A.B.C.D
IPv4 address of the TWAMP Reflector

HOSTNAME
The hostname of the TWAMP Reflector

Default

Mandatory attribute

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1.

Example

The following example shows the Reflector IP being configured in the TWAMP Session PE1-SESSION.

```
OcNOS#configure terminal
OcNOS(config)# twamp session PE1-SESSION
OcNOS(twamp-session)# reflector-ip 203.1.1.1
```

© 2026 IP Infusion Inc. Proprietary

reflector-port

Use this command to set the UDP port of the TWAMP Reflector so that this TWAMP Session will send the test packets.

Command Syntax

```
reflector-port (862|<1025-65535>)
```

Parameters

862
Default UDP port for TWAMP

<1025-65535>
The UDP port to be configured

Default

862

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1.

Example

The following example shows the reflector-port UDP port being configured in the TWAMP Session PE1-SESSION.

```
OcNOS#configure terminal
OcNOS(config)# twamp session PE1-SESSION
OcNOS(twamp-session)# reflector-port 862
```

© 2026 IP Infusion Inc. Proprietary

sender-ip

Use this command to set the IP of the TWAMP Sender so that this TWAMP Session will send the test packets. The TWAMP Sender is expected to be the IP from the Access Interface in which the TWAMP test packets are going to be sent.

Command Syntax

```
sender-ip (A.B.C.D|HOSTNAME)
```

Parameters

A.B.C.D
IPv4 address of the TWAMP Sender

HOSTNAME
The hostname of the TWAMP Sender

Default

Mandatory attribute

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1.

Example

The following example shows the Sender IP being configured in the TWAMP Session PE1-SESSION.

```
OcNOS#configure terminal
OcNOS(config)# twamp session PE1-SESSION
OcNOS(twamp-session)# sender-ip 101.1.1.1
```

© 2026 IP Infusion Inc. Proprietary

sender-port

Use this command to set the UDP port of the TWAMP Sender so that this TWAMP Session will send the test packets.

Command Syntax

```
sender-port (862|<1025-65535>)
```

Parameters

862
Default UDP port for TWAMP

<1025-65535>
The UDP port to be configured

Default

862

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1

Example

The following example shows the sender-port UDP port being configured in the TWAMP Session PE1-SESSION.

```
OcNOS#configure terminal
OcNOS(config)# twamp session PE1-SESSION
OcNOS(twamp-session)# sender-port 862
```

© 2026 IP Infusion Inc. Proprietary

dscp

Use this command to set the DSCP (Differentiated Services Code Point) of the TWAMP test packets.

Command Syntax

```
dscp WORD
```

Parameters

WORD

Value can be 0-63 or type well-known names (af11| af12| af13| af21| af22| af23| af31| af32| af33| af41| af42| af43| cs1| cs2| cs3| cs4| cs5| cs6| cs7| default| ef).

Default

Default

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1.

Example

The following example shows the dscp being configured in the TWAMP Session PE1-SESSION.

```
OcNOS#configure terminal
OcNOS(config)# twamp session PE1-SESSION
OcNOS(twamp-session)# dscp af23
```

© 2026 IP Infusion Inc. Proprietary

vrf-name

Use this command to set the VRF to be used when sending test packets over the TWAMP Reflector IP, which is the Access interface.

Command Syntax

```
vrf-name WORD
```

Parameters

WORD
VRF name

Default

default

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1

Example

The following example shows the VRF being configured, assigned to interface xe4 and being configured in the TWAMP Session PE1-SESSION.

```
OcNOS#configure terminal
OcNOS(config)# ip vrf vrf101
OcNOS(config-if)# interface xe4.101
OcNOS(config-if)# ip vrf forwarding vrf101
OcNOS(config-if)# ip address 103.1.1.1/24
OcNOS(config-if)# exit
OcNOS(config)# twamp session PE1-SESSION
OcNOS(twamp-session)# vrf-name vrf101
OcNOS(twamp-session)# sender-ip 103.1.1.1
```

© 2026 IP Infusion Inc. Proprietary

vpn-net-interface

When this attribute is configured, the TWAMP session will execute the inter-access measurement. Use this command to configure the physical network interface from which the test packets should go out.

Command Syntax

```
vpn-net-interface WORD
```

Parameters

WORD
Interface name

Default

Not applicable

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1

Example

The following example shows the vpn-net-interface being configured on TWAMP Session PE1-SESSION.

```
OcNOS#configure terminal
OcNOS(config)# ip vrf vrf101
OcNOS(config-if)# interface xe4.101
OcNOS(config-if)# ip vrf forwarding vrf101
OcNOS(config-if)# ip address 103.1.1.1/24
OcNOS(config-if)# exit
OcNOS(config)# twamp session PE1-SESSION
OcNOS(twamp-session)# vrf-name vrf101
OcNOS(twamp-session)# sender-ip 103.1.1.1
OcNOS(twamp-session)# vpn-net-interface xe7
```

© 2026 IP Infusion Inc. Proprietary

description

Use this command to configure a description in the TWAMP Session.

Command Syntax

```
description WORD
```

Parameters

WORD
Description to be added

Default

None

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1

Example

The following example shows the description being set in TWAMP Session PE1-SESSION.

```
OcNOS#configure terminal
OcNOS(config)# twamp session PE1-SESSION
OcNOS(twamp-session)# description TWAMP VPN Inter-Acc-Intf
```

© 2026 IP Infusion Inc. Proprietary

burst-interval

Use this command to configure the burst interval between test packets sent to TWAMP Reflector.

Command Syntax

```
burst-interval <1000-15000>
```

Parameters

<1000-15000>
The burst interval (in seconds) between test packets

Default

3000

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1

Example

The following example shows how to configure the burst-interval in TWAMP Session PE1-SESSION.

```
OcNOS#configure terminal
OcNOS(config)# twamp session PE1-SESSION
OcNOS(twamp-session)# burst-interval 10000
```

© 2026 IP Infusion Inc. Proprietary

burst-count

Use this command to configure the number of test packets to be sent in a burst.

Command Syntax

```
burst-count <1-10>
```

Parameters

<1-10>

Set the number of packets to be sent at each burst interval. The default value is 3.

Note: The limit **<1-10>** is applicable only on Qumran2 (Q2) series platforms. For Qumran1 (Q1) series platforms it is llimited to **<1-5>**.

Default

3

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1. The burst count limit has been increased from on <1-5> to <1-10> on Qumran2 (Q2) series platformsin OcNOS version 7.0.0.

Example

The following example shows how to configure the burst-interval in TWAMP Session PE1-SESSION.

```
OcNOS#configure terminal
OcNOS(config)# twamp session PE1-SESSION
OcNOS(twamp-session)# burst-count 4
```

© 2026 IP Infusion Inc. Proprietary

interval

Use this command to configure the computation interval of test packets. The computation interval is the interval where the measurement statistics are calculated periodically.

Command Syntax

```
interval <30-3600>
```

Parameters

<30-3600>
The computation interval where the measurement statistics are calculated.

Default

3

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1.

Example

The following example shows how to configure the burst-interval in TWAMP Session PE1-SESSION.

```
OcNOS#configure terminal
OcNOS(config)# twamp session PE1-SESSION
OcNOS(twamp-session)# burst-count 4
```

© 2026 IP Infusion Inc. Proprietary

notification enabled

Use this command to enable the generation of notifications for TWAMP Session when the average delay crosses the configured threshold or any packet-loss is seen. The generated notification for each event is sent over NETCONF Notification and SNMP Trap.

For threshold crossing, please check the other notification commands: "`notification threshold`" and "`notification minimum-change`" for more details on how to finely adjust the generation of the notifications.

For the packet-loss, there is no fine adjustment. Once a packet loss is detected, the notification is generated.

Command Syntax

```
notification enabled
```

Parameters

None

Default

Disabled

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1.

Example

The following example shows how to configure the notification in TWAMP Session PE1-SESSION.

```
OcNOS(config)# twamp session PE1-SESSION
OcNOS(twamp-session)# notification enabled
```

© 2026 IP Infusion Inc. Proprietary

notification threshold

Use this command to configure the notification threshold, which will trigger the notification in case the average delay crosses the configured threshold. The threshold is a percentage value with respect to the average delay. Thus, if the average delay changes above or below the configured threshold, then the NETCONF Notification and SNMP Trap are generated to inform that the threshold has been crossed.

Note: If notification minimum-change is used together with the notification threshold, then the one that crosses the threshold first will trigger the NETCONF Notification and SNMP Trap.

Command Syntax

```
notification threshold <1-100>
```

Parameters

<1-100>
The percentile threshold with respect to the average delay calculated

Default

30

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1.

Example

The following example shows how to configure the notification in TWAMP Session PE1-SESSION.

```
OcNOS#configure terminal
OcNOS(config)# twamp session PE1-SESSION
OcNOS(twamp-session)# notification enabled
OcNOS(twamp-session)# notification threshold 80
```

© 2026 IP Infusion Inc. Proprietary

notification minimum-change

Use this command to configure the notification minimum-change threshold. The minimum-change threshold unit is in microseconds and is related to the average delay. Thus, if the average delay changes above or below the minimum-change threshold, then the NETCONF Notification and SNMP Trap are generated to inform that the threshold has been crossed.

Observation: If notification minimum-change is used together with the notification threshold, then the one that crosses the threshold first will trigger the NETCONF Notification and SNMP Trap.

Command Syntax

```
notification minimum-change <0-10000>
```

Parameters

<1-10000>
The threshold in microseconds with respect to the average delay calculated

Default

5000

Command Mode

TWAMP-SESSION mode

Applicability

This command is introduced in OcNOS version 6.6.1.

Example

The following example shows how to configure the notification in TWAMP Session PE1-SESSION.

```
OcNOS#configure terminal
OcNOS(config)# twamp session PE1-SESSION
OcNOS(twamp-session)# notification enabled
OcNOS(twamp-session)# notification threshold 80
OcNOS(twamp-session)# notification minimum-change 80
```

© 2026 IP Infusion Inc. Proprietary

show twamp-sessions

Use this command to show a summary of the last calculated statistics of the running TWAMP Sessions. These statistics are related to the computation interval configured.

Command Syntax

```
show twamp-sessions
```

Parameters

None

Default

None

Command Mode

Execution modeand Privilege mode

Applicability

This command is introduced in OcNOS version 6.6.1.

Example

The following example shows the TWAMP sessions summary .

```
#show twamp-session
Session-Name  Net-Intf Delay(us) Min(us) Max(us) Var(us) Loss(%)   State
PE1-SESSION     xe7      22        20      23      0       0.00    Running
```

Table 106. Table 1-1: show twamp-sessions fields

Field	Description
Session-Name	TWAMP Session session name
Net-Intf	Network Interface where this TWAMP Session is
Delay(us)	Average Round-Trip Delay
Min(us)	Average Minimum Round-Trip Delay
Max(us)	Average Maximum Round-Trip-Delay
Var(us)	Average Round-Trip Delay Variation
Loss(%)	Percentage of packet-loss

© 2026 IP Infusion Inc. Proprietary

show twamp-sessions WORD

Use this command to show detailed statistics of the running TWAMP Session. These statistics are related to the computation interval configured.

Command Syntax

```
show twamp-sessions WORD
```

Parameters

WORD
The TWAMP Session session name

Default

Not Applicable

Command Mode

Execution mode and Privileged Execution mode

Applicability

This command is introduced in OcNOS version 6.6.1.

Example

The following example shows the detailed statistics of TWAMP session.

```
#show twamp-session PE1-SESSION
Session Name      : PE1-SESSION
VPN Network Interface name    : xe7
State             : Running
Sender IP         : 103.1.1.1
Sender Port       : 862
DSCP value        : default
Vrf-Name          : vrf101
Reflector IP      : 203.1.1.1
Reflector port    : 862
Execution Status  : Running
Last Calculated stats:
Time: 2025-06-12 12:04:22
  Average delay          : 22
  Minimum delay          : 20
  Maximum delay          : 23
  Average delay variation: 0
  Minimum delay variation: 0
  Maximum delay variation: 0
  Packets sent           : 30
  Packets received       : 30
  Packets timeout        : 0
Packet Loss : 0.00
Notification stats:
  Notification : Enabled
  Threshold (%): 80
  Minimum-change (usec): 80
  Last Notified Packet Loss:
    Time: 2025-06-12 12:12:12
    Packet Loss: 30.00
  Last Notified Average Delay:
```

```
    Time: -
    Average Delay (usec):
```

Table 107. Table 1-1: show twamp-sessions WORD fields

Field	Description
VPN Network Interface name	TAMP Session session name
Sender IP	Network Interface where this TWAMP Session is
Sender Port	Average Round-Trip Delay
DSCP Value	Average Minimum Round-Trip Delay
Vrf-Name	Average Maximum Round-Trip-Delay
Reflector IP	Average Round-Trip Delay Variation
Reflector Port	Percentage of packet-loss
Execution status	Shows if session is running or not
HW Status	Indicate if Hardware filters are configured
Last Calculated stats	Statistics related to the last calculated values
Time	When the last calculation was done. This is related to the calculation interval configuration.
Average delay	Average delay last calculated.
Minimum delay	Minimum delay last calculated.
Maximum delay	Maximum delay last calculated.
Average delay variation	Average delay variation last calculated.
Minimum delay variation	Minimum delay variation last calculated.
Maximum delay variation	Maximum delay variation last calculated.
Packets sent	How many packets were sent
Packets received	How many packets were received
Packets	How many packets were timed out

© 2026 IP Infusion Inc. Proprietary

Table 107. show twamp-sessions WORD fields (continued)

Field	Description
timeout	
Packet Loss	Percentage of packet loss
Notification stats	Statistics related to the notification values
Notification	Notification is enabled or not
Threshold	Configured threshold
Minimum-change	Configured minimum-change
Average delay	Average delay to be considered for the thresholds. This value is set when we enable the “notification enabled” command.
Last Notified Packet Loss	Statistics of when last notified packet loss was sent
Time	The time when last notified packet loss was sent
Packet Loss	What was the value notified
Last Notified Average Delay	Statistics of when last notified average delay was sent
Time	The time when last notified average delay was sent
Average Delay	What was the value notified

© 2026 IP Infusion Inc. Proprietary

show twamp-sessions all

Use this command to show detailed statistics of all the running TWAMP Sessions. These statistics are related to the computation interval configured.

Command Syntax

```
show twamp-sessions all
```

Parameters

None

Default

None

Command Mode

Execution mode and Privileged execution mode

Applicability

This command is introduced in OcNOS version 6.6.1.

Example

The following example shows the detailed statistics of all the TWAMP sessions configured .

```
OcNOS#configure terminal
PE1#sh twamp-sessions all
Session Name      : PE1-SESSION
VPN Network Interface name    : xe7
State             : Running
Sender IP         : 103.1.1.1
Sender Port       : 862
DSCP value        : default
Vrf-Name          : vrf101
Reflector IP      : 203.1.1.1
Reflector port    : 862
Execution Status  : Running
Last Calculated stats:
Time: 2025-06-12 12:04:22
  Average delay          : 22
  Minimum delay          : 20
  Maximum delay          : 23
  Average delay variation: 0
  Minimum delay variation: 0
  Maximum delay variation: 0
  Packets sent           : 30
  Packets received       : 30
  Packets timeout        : 0
Packet Loss : 0.00
Notification stats:
  Notification : Enabled
  Threshold (%): 80
  Minimum-change (usec): 80
  Last Notified Packet Loss:
    Time: 2025-06-12 12:12:12
    Packet Loss: 30.00
  Last Notified Average Delay:
    Time: -
```

```
    Average Delay (usec): -
```

Table 108. Table 1-1: show twamp-sessions WORD fields

Field	Description
Session Name	TWAMP Session session name
VPN Network Interface name	VPN Network interface name
Start Time	When the TWAMP Session started
Elapsed time (milli sec)	Elapsed time
Packets Sent	How many packets sent
Packets Received	How many packets received
Packet Loss (%)	Percentage of packet loss
Round Trip Delay(usec)	Round Trip Delay calculations
Minimum	Round Trip Delay Minimum
Maximum	Round Trip Delay Maximum
Average	Round Trip Delay Average
Forward Delay(usec)	Forward Delay calculations
Minimum	Forward Delay Minimum
Maximum	Forward Delay Maximum
Average	Forward Delay Average
Reverse Delay(usec)	Reverse Delay calculations
Minimum	Reverse Delay Minimum
Maximum	Reverse Delay Maximum
Average	Reverse Delay Average
Round Trip Delay Variation (usec)	Round Trip Delay Variation calculations
Minimum	Round Trip Delay Variation Minimum
Maximum	Round Trip Delay Variation Maximum
Average	Round Trip Delay Variation Average

© 2026 IP Infusion Inc. Proprietary

Table 108. show twamp-sessions WORD fields (continued)

Field	Description
Forward Delay Variation (usec)	Forward Delay Variation calculations
Minimum	Forward Delay Variation Minimum
Maximum	Forward Delay Variation Maximum
Average	Forward Delay Variation Average
Reverse Delay Variation (usec)	Reverse Delay Variation calculations
Minimum	Reverse Delay Variation Minimum
Maximum	Reverse Delay Variation Maximum
Average	Reverse Delay Variation Average

© 2026 IP Infusion Inc. Proprietary

show twamp packet-rate

Use this command to display the overall TWAMP packet transmission rate in packets per second (pps) for both client and server operations.

Command Syntax

```
show twamp  packet-rate
```

Parameters

None

Default

None

Command Mode

Execution modeand Privilege mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

The following example shows the TWAMP packet-rate summary .

```
ocnos#show twamp  packet-rate
! [execution timestamp : 2025 Oct 24 13:42:04]
TWAMP packet sending rate: 2996 pps
             senders rate: 2996 pps
          reflectors rate: 0 pps
Maximum supported sending rate: 5120 pps
```

Table 109. Table 1-1: show twamp packet-rate fields

Field	Description
TWAMP packet sending rate	Shows the total number of TWAMP test packets transmitted per second (pps) by all active TWAMP sessions on the device.
senders rate	Indicates the cumulative packet sending rate from all TWAMP clients (senders) currently transmitting measurement packets.
reflectors rate	Displays the packet response rate from TWAMP reflectors (servers). A value of 0 pps means no active reflector sessions are responding at this time.
Maximum supported sending rate	Defines the maximum TWAMP packet transmission rate (in pps) that the device can support. This limit ensures session and traffic control within hardware or software capacity.

© 2026 IP Infusion Inc. Proprietary

show twamp session count

Use this command to display the total number of active TWAMP sessions that are currently established.

Command Syntax

```
show twamp session-count
```

Parameters

None

Default

None

Command Mode

Execution modeand Privilege mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

The following example shows the TWAMP session summary .

```
ocnos#show twamp session-count
! [execution timestamp : 2025 Oct 24 13:26:31]
TWAMP session count:
  Link measurement sessions: 500

  End to End sessions      : 1000
    Running          : 1000
    Not Running      : 0

  Reflector sessions       : 0

  Configure Client sessions: 500
    Idle             : 499
    Accepted         : 0
    Not Accepted     : 0
    Started          : 0
    Stopped          : 0

  Accepted Server sessions : 0

TOTAL                      : 2000 sessions
```

Table 110. Table 1-1: show twamp session-countfields

Field	Description
Link measurement sessions	Number of TWAMP sessions used for measuring link-level performance.
End to End sessions	Number of TWAMP sessions configured for end-to-end performance measurement (multi-hop tests).

Table 110. show twamp session-countfields (continued)

Field	Description
Running	Number of end-to-end TWAMP sessions that are currently active and transmitting test packets.
Not Running	Number of end-to-end TWAMP sessions that are configured but not currently active.
Reflector sessions	Number of TWAMP reflector (server-side) sessions currently active on the device.
Configure Client sessions	Number of TWAMP client sessions configured on the device.
Idle	Number of configured client sessions that are inactive or waiting for initiation.
Accepted	Number of TWAMP client sessions that have been accepted by the server.
Not Accepted	Number of TWAMP client sessions that were not accepted by the server (possibly due to configuration mismatch or policy restrictions).
Started	Number of TWAMP client sessions that have been successfully started and are transmitting traffic.
Stopped	Number of TWAMP client sessions that were previously active but have been stopped.
Accepted Server sessions	Number of TWAMP sessions accepted by the device when acting as a TWAMP server.
TOTAL	The total number of TWAMP sessions, summing up all session categories (link measurement + end-to-end + client/server).

© 2026 IP Infusion Inc. Proprietary

Glossary

CE	Customer Edge
PE	Provider Edge
VPN	Virtual Private Network

© 2026 IP Infusion Inc. Proprietary

TWO-WAY ACTIVE MEASUREMENT PROTOCOL COMMAND REFERENCE

TWAMP Commands 2840
control-admin-state 2842
reflector-admin-state 2843
twamp stop-test-session (Deprecated) 2844
reflector-name 2845
show debugging twamp 2847
show running-config twamp 2848
Show twamp server connections information 2849
show twamp server sessions information 2850
show twamp server status 2851
show twamp statistics 2852
twamp client 2854
twamp client reset connection 2855
twamp client start connection 2856
twamp client stop connection 2857
twamp connection 2858
twamp server 2859
twamp server inactivity-timeout 2860
twamp server listen-vrf 2861
twamp server max-connection-duration 2862
twamp server max-connections 2863
twamp server maximum-sessions-per-connection 2864
twamp client maximum-connections 2865
twamp client maximum-sessions 2866
twamp client maximum-sessions-per-connection 2867
twamp client server 2868
twamp client session 2869
twamp-light control 2870
twamp-light reflector 2871
twamp server admin state 2872
twamp server max-connections-per-client 2873
twamp server maximum-sessions 2874
twamp server port 2875
twamp start-test-session (Deprecated) 2876
test-session-name (Deprecated) 2877

© 2026 IP Infusion Inc. Proprietary

TWAMP Commands

This section describes the commands used to manage the Two-Way Active Measurement Protocol (TWAMP).

Notes:

- The implementation of TWAMP measurement relies on software timestamps collected at the application layer. The measured packet path in this case is longer than the actual link path and thus the timestamps used in the delay measurement are less accurate.
- On Qumran-UX (QUX) series platforms, where resources are limited, activate TWAMP hardware profiles only when necessary, as other profiles may require the same board resources. For more details, refer to the "`hardware-profile filter`" commands in the **System Configure Mode Commands** section of the *OcNOS System Management document*.

This section describes the commands for TWAMP.

control-admin-state 2842
reflector-admin-state 2843
twamp stop-test-session (Deprecated) 2844
reflector-name 2845
show debugging twamp 2847
show running-config twamp 2848
Show twamp server connections information 2849
show twamp server sessions information 2850
show twamp server status 2851
show twamp statistics 2852
twamp client 2854
twamp client reset connection 2855
twamp client start connection 2856
twamp client stop connection 2857
twamp connection 2858
twamp server 2859
twamp server inactivity-timeout 2860
twamp server listen-vrf 2861
twamp server max-connection-duration 2862
twamp server max-connections 2863
twamp server maximum-sessions-per-connection 2864
twamp client maximum-connections 2865
twamp client maximum-sessions 2866
twamp client maximum-sessions-per-connection 2867
twamp client server 2868

© 2026 IP Infusion Inc. Proprietary

twamp client session 2869
twamp-light control 2870
twamp-light reflector 2871
twamp server admin state 2872
twamp server max-connections-per-client 2873
twamp server maximum-sessions 2874
twamp server port 2875
twamp start-test-session (Deprecated) 2876
test-session-name (Deprecated) 2877

© 2026 IP Infusion Inc. Proprietary

control-admin-state

Use this command to enable or disable the TWAMP sender framework. The sender framework needs to be enabled before configuring any TWAMP sessions.

Command Syntax

```
control-admin-state (enable | disable)
```

Parameters

enable
Enable the TWAMP sender framework.

disable
Disable the TWAMP sender framework.

Default

Disable

Command Mode

TWAMP light control mode

Applicability

This command was introduced before OcNOS version 3.0.

Example

```
(config)#twamp-light control
(config-twamp-light-ctrl)#control-admin-state enable
```

© 2026 IP Infusion Inc. Proprietary

reflector-admin-state

Use this command to enable or disable the TWAMP reflector framework. The reflector framework needs to be enabled before configuring any reflector session.

Command Syntax

```
reflector-admin-state (enable | disable)
```

Parameters

enable
Enable the TWAMP reflector framework.

disable
Disable the TWAMP reflector framework.

Default

Disabled

Command Mode

TWAMP light reflector mode (config-twamp-light-ctrl)

Applicability

This command was introduced before OcNOS version 3.0.

Example

```
(config)#twamp-light reflector
(config-twamp-light-ctrl)#reflector-admin-state enable
```

© 2026 IP Infusion Inc. Proprietary

twamp stop-test-session (Deprecated)

This command is not supported anymore. The “delay-measurement dynamic twamp” command under interface mode should be used to start/stop twamp measurement instead.

Use this command to stop the TWAMP test session.

Command Syntax

```
twamp stop test-session WORD
```

Parameters

WORD

TWAMP test session name.

Default

None

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 3.0.

Examples

```
#twamp stop test-session s1
```

© 2026 IP Infusion Inc. Proprietary

reflector-name

Use this command to configure a TWAMP light reflector session.

Command Syntax

```
reflector-name WORD reflector-ip (ipv4 A.B.C.D|ipv6 X:X::X:X|HOSTNAME) (reflector-port (862|<1025-
65535>)|) (vrf NAME|) (sender-ip (ipv4 A.B.C.D|ipv6 X:X::X:X|HOSTNAME)|) (sender-port (862|<1025-
65535>)|)(vpn-net-interface IFNAME|)
```

Parameters

WORD
Set TWAMP Light reflector session name

X:X::X:X
TWAMP Reflector IP address in IPv6 format.

A.B.C.D
TWAMP Reflector IP address in IPv4 format.

HOSTNAME
The hostname of the TWAMP reflector

vrf NAME
Set TWAMP Light reflector session VRF name

reflector-port
Set the TWAMP Reflector UDP port number where the reflector will be listening

sender-ip
Set the sender IP from which the reflector will accept the TWAMP packets. If the source of the TWAMP packet does not match this address, the packet will be ignored. In case not set, the reflector will not discriminate on the sender IP value.

sender-port
Set the sender UDP port from which the reflector will accept the TWAMP packets. In case not set, the reflector will not discriminate on the sender UDP port value.

vpn-net-interface
Set the VPN Network Interface from which the TWAMP test packets must go out. This parameter is needed only in case of VPN Inter-Access Interface.

Default

862 for reflector-port

0 for the sender-port (interpreted as no filtering on sender port)

Command Mode

TWAMP light reflector mode (config-twamp-light-ctrl)

Applicability

This command was introduced before OcNOS version 5.1.

© 2026 IP Infusion Inc. Proprietary

The sender-IP and sender-port parameters were added in OcNOS version 6.0.0

Example

```
(config)#twamp-light reflector
(config-twamp-light-ref)# reflector-name R11 reflector-ip 23.1.1.1 vrf vrf1
(config-twamp-light-ref)# reflector-name R12 reflector-ip 24.1.1.1 vrf vrf2
(config-twamp-light-ref)# reflector-name R12 reflector-ip 24.1.1.1 sender-ip 24.1.1.2 vrf vrf2
(config-twamp-light-ref)# reflector-name R12 reflector-ip 24.1.1.1 sender-ip 24.1.1.2 sender-port
1555 vrf vrf2
(config-twamp-light-ref)# reflector-name R15 reflector-ip 10.1.11.1 vrf vrf2 vpn-net-interface xe4
```

© 2026 IP Infusion Inc. Proprietary

show debugging twamp

Use this command to display the status of debug information of the TWAMP.

Command Syntax

```
show debugging twamp
```

Parameters

None

Default

Off

Config Mode

Execution mode and Privilege mode

Applicability

This command was introduced in OcNOS version 6.1.0.

Examples

```
OcNOS#show debugging twamp
TWAMP debugging status: OFF
```

show running-config twamp

Use this command to display TWAMP running configuration alone.

Command Syntax

```
show running-config twamp
```

Parameters

None

Command Mode

Execution modeand Privileged execution mode

Applicability

This command was introduced before OcNOS version 3.0.

Examples

Reflector:

```
#show running-config twamp
twamp-light reflector
reflector-admin-state enable
reflector-ip 10.0.0.1 reflector-port 1234
```

Control or Sender:

```
#show running-config twamp
twamp-light control
 control-admin-state enable
 test-session-name s1 sender-ip 10.0.0.1 sender-port 1234 reflector-ip 10.0.0.1 reflector-port 1234
```

© 2026 IP Infusion Inc. Proprietary

Show twamp server connections information

Show the existing connections information.

Command Syntax

```
show twamp server connection (WORD|)
```

Parameters

WORD
The connection ID

Default

None

Command mode

Execution mode

Applicability

This command was introduced in OcNOS version 6.2.0

Examples

```
OcNOS#  show twamp server connectionConnection ID Client Address   Client
Port            1      100.1.0.1         59807           2      100.1.0.2         59808           3
      100.1.0.3         59809           4      100.1.0.4         59810
OcNOS#  show twamp server connection 2Connection ID              : 2Client Address             :
100.1.0.2Client Port               : 59808Client VRF                 :
defaultState                     : Control session establishedConnection Time           : Thu Feb
14 12:13:51 2019Last Activity Time         : Thu Feb 14 12:13:51 2019Number of session requests :
7Number of accepted sessions: 3
OcNOS#
```

show twamp server sessions information

Show the existing sessions information

Command Syntax

```
show twamp server session (WORD|)
```

Parameters

WORD
The session ID

Default

None

Command mode

Execution mode

Applicability

This command was introduced in OcNOS version 6.2.0

Examples

```
OcNOS#  show twamp server session
Connection ID                           Session
Id                                  Sender                                  Receiver       State
            1
5B.5B.5B.03:E00FD9FF99B7F173:D8CE5E5D                                91.91.91.2
        91.91.91.3     Started
            2
5B.5B.5B.03:E00FCF8221DF548E:EF824500                                   9.1.9.1
        91.91.91.3     Started

OcNOS# show twamp server session 5B.5B.5B.03:E00FCF8221DF548E:EF824500
State           : Started
Connection ID   : 2
Session ID      : 5B.5B.5B.03:E00FCF8221DF548E:EF824500
Sender address  : 9.1.9.1
Sender port     : 44566
Receiver address: 91.91.91.3
Receiver port   : 44001
Receiver VRF    : default
Start time      : 2022-08-31 16:44:00
Timeout         : 5 sec
OcNOS#
```

© 2026 IP Infusion Inc. Proprietary

show twamp server status

Show the general status of the TWAMP server

Command Syntax

```
show twamp server status
```

Parameters

None

Default

None

Command mode

Exec mode

Applicability

This command was introduced in OcNOS version 6.2.0.

Examples

```
OcNOS#show twamp server statusStatus                    : EnabledListening Port            :
862Server VRF list           : management default vrf1Number of connections     : 3Number of
sessions        : 15Number of running sessions: 5
OcNOS#
```

© 2026 IP Infusion Inc. Proprietary

show twamp statistics

Use this command to display TWAMP test session.

Command Syntax

```
show twamp-statistics (WORD|)
```

Parameters

WORD
TWAMP test session name.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 3.0.

Examples

```
#show twamp statistics
=======================================
      TWAMP Test-Session Statistics
=======================================
  Test-Session Name        : s1
  Start Time               : 2019 Aug 05 08:40:12
  Elapsed time(milli sec)  : 4004
  Packets Sent             : 4
  Packets Received         : 4
  Packet Loss(%)           : 0.0000
  Round Trip Delay(usec)
      Minimum              : 585
      Maximum              : 2306
      Average              : 1038
  Forward Delay(usec)
      Minimum              : 254
      Maximum              : 318
      Average              : 280
  Reverse Delay(usec)
      Minimum              : 331
      Maximum              : 1988
      Average              : 758
  Round Trip Delay Variation(usec)
      Minimum              : 2306
      Maximum              : 2640
      Average              : 2444
  Forward Delay Variation(usec)
      Minimum              : 318
      Maximum              : 375
      Average              : 345
  Reverse Delay Variation(usec)
      Minimum              : 1988
      Maximum              : 2265
      Average              : 2099
```

© 2026 IP Infusion Inc. Proprietary

Table 111. show twamp statistics fields

Field	Description
Test-Session Name	TWAMP test session name
Start Time	TWAMP test start time
Elapsed time (milli sec)	Time taken to complete the test in milliseconds
Packets Sent	Number of packet sent
Packets Received	Number of packet received
Packet Loss (%)	Packet lost in percentage
Round Trip Delay(usec)	Round trip delay between TWAMP sender and reflector. Displayed Minimum, Maximum and Average round trip delay in micro-seconds
Forward Delay(usec)	Forward delay between TWAMP sender and reflector. Displayed Minimum, Maximum and Average forward delay in micro-seconds In case the clocks on the sender and reflector are not in sync, a (*) will be displayed here.
Reverse Delay(usec)	Reverse delay between TWAMP reflector and sender. Displayed Minimum, Maximum and Average reverse delay in micro-seconds In case the clocks on the sender and reflector are not in sync, a (*) will be displayed here.
Round Trip Delay Variation (usec)	Round trip delay variation between TWAMP sender and reflector. Displayed Minimum, Maximum and Average round trip delay variation in micro-seconds
Forward Delay Variation (usec)	Forward delay variation between TWAMP sender and reflector. Displayed Minimum, Maximum and Average forward delay variation in micro-seconds In case the clocks on the sender and reflector are not in sync, a (*) will be displayed here.
Reverse Delay Variation (usec)	Reverse delay variation between TWAMP reflector and sender. Displayed Minimum, Maximum and Average reverse delay variation in micro-seconds In case the clocks on the sender and reflector are not in sync, a (*) will be displayed here.

© 2026 IP Infusion Inc. Proprietary

twamp client

Use this command to enable and configure the TWAMP client.

Command Syntax

```
twamp client
```

Parameters

None

Default

None

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
OcNOS(config)#
OcNOS(config)#twamp client
OcNOS(config-twamp-client)#
```

twamp client reset connection

Use this command to reset the TWAMP client sessions on a connection.

Command Syntax

```
twamp client reset connection NAME (session NAME|) (vrf WORD|)
```

Parameters

NAME
Specifies the name of the connection where the sessions will be reset.

NAME
Specifies the name of the session to be reset. If not specified, all sessions on the connection will be reset.

WORD
Specifies the name of the vrf where the connection is operating. If not specified, the default vrf will be used.

Default

None

Command mode

Execution mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
OcNOS#twamp client reset connection test_ipv6 vrf vrf1
```

© 2026 IP Infusion Inc. Proprietary

twamp client start connection

Use this command to start the sessions already requested from the server and accepted by the server over the specified connection.

Command Syntax

```
twamp client start connection (NAME) packet-count (0-1000000) (vrf NAME)
```

Parameters

NAME
Specifies the name of the TWAMP client connection to be used for the session.

0-1000000
Specifies the number of packets to be used in the session. It can be any value between 0 and 1,000,000.

vrf NAME
Specifies the a Virtual Routing and Forwarding (VRF) context for the TWAMP client session (optional).

Default

None

Command mode

exec mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
OcNOS(config)#twamp client
OcNOS(config-twamp-client)#twamp client start connection MyConnection packet-count 1000 vrf VRF1
OcNOS(config-twamp-client)#commit
```

© 2026 IP Infusion Inc. Proprietary

twamp client stop connection

Use this command to stop the sessions already requested from the server and accepted by the server over the specified connection.

This command is only effective if the packet-count specified in the start command is not elapsed. In case the packet-count is elapsed, the sessions are stopped automatically, and this command would not have any effect.

Command Syntax

```
twamp client stop connection (NAME)(vrf NAME)
```

Parameters

NAME
Specifies the name of the TWAMP client connection to be stopped.

vrf NAME
Specifies the a VRFcontext for the TWAMP client session (optional).

Default

None

Command mode

Execution mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
OcNOS(config)#twamp client
OcNOS(config-twamp-client)#twamp client stop connection MyConnection vrf VRF1
OcNOS(config-twamp-client)#commit
```

© 2026 IP Infusion Inc. Proprietary

twamp connection

Use this command to create a connection between the server and the client.

Multiple connections can be specified. The user can issue the same command multiple times, once for each VRF.

Only clients in the configured VRFs are allowed to connect to the server and only sessions to IPs in the configured VRFs are allowed to be established.

Use the no form of this command to remove the connection.

Command Syntax

```
connection NAME (vrf-name NAME)
no connection NAME (vrf-name NAME)
```

Parameters

NAME

Specifies the connection name to add to the client.

NAME

Specifies the VRF name.

Default

None

Command mode

twamp-client mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
OcNOS(config)#twamp client
OcNOS(config-twamp-client)#connection Connection1 vrf VRF1
OcNOS(config-twamp-client)#commit
OcNOS(config-twamp-client)#
OcNOS(config-twamp-client)#no connection Connection1 vrf VRF1
OcNOS(config-twamp-client)#commit
```

© 2026 IP Infusion Inc. Proprietary

twamp server

Use this command to change to the twamp server mode.

Command Syntax

```
twamp server
```

Parameters

None

Default

None

Command mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.2.0

Examples

```
OcNOS(config)#
OcNOS(config)#twamp server
OcNOS(config-twamp-server)#
```

twamp server inactivity-timeout

Use this command to set the inactivity timeout in minutes. In case the server does not receive any control packet from the client during this timeout, the connection to the client is dropped and any session started by this client will be stopped.

The default value of this flag is 0. In this case the server does not check for inactivity from the client.

Use the no form of this command to set the value to the default value 0.

Command Syntax

```
inactivity-timeout <0-30>
```

Parameters

<0-30>
Set the inactivity timeout to a value between 0 and 30 in minutes.

Default

Zero

Command mode

twamp-server mode

Applicability

This command was introduced in OcNOS version 6.2.0

Examples

```
OcNOS(config)#twamp server
OcNOS(config-twamp-server)#inactivity-timeout 22
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#no inactivity-timeout
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#
```

© 2026 IP Infusion Inc. Proprietary

twamp server listen-vrf

Use this command to a dd custom VRFs to the list of VRFs where the TWAMP server is running. If no listen-vrf is specified, the TWAMP server is started in the default VRF only.

Multiple VRFs can be specified. The user can issue the same command multiple times, once for each VRF.

Only clients in the configured VRFs are allowed to connect to the server and only sessions to IPs in the configured VRFs are allowed to be established.

Use the no form of this command to delete a listen-vrf.

Command Syntax

```
listen-vrf <NAME>
no listen-vrf <NAME>
```

Parameters

NAME

VRF name to add to the server.

Default

Default

Command mode

twamp-server mode

Applicability

This command was introduced in OcNOS version 6.2.0.

Examples

```
OcNOS(config)#twamp server
OcNOS(config-twamp-server)#listen-vrf default
OcNOS(config-twamp-server)#listen-vrf management
OcNOS(config-twamp-server)#listen-vrf vrf1
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#
OcNOS(config-twamp-server)#no listen-vrf management
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#
```

© 2026 IP Infusion Inc. Proprietary

twamp server max-connection-duration

Use this command to set the maximum connection duration that a client can establish with the server in hours. The default value of this flag is 0, which means that there is no limit for the connection duration.

Use the no form of this command to set the value to the default value 0.

Command Syntax

```
max-connection-duration <0-120>
```

Parameters

<0-120>
Set the maximum connection duration value between 0 and 120 hours.

Default

Zero

Command mode

twamp-server mode

Applicability

This command was introduced in OcNOS version 6.2.0

Examples

```
OcNOS(config)#twamp server
OcNOS(config-twamp-server)#max-connection-duration 5
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#no max-connection-duration
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#
```

© 2026 IP Infusion Inc. Proprietary

twamp server max-connections

Use this command to set the maximum number of connections that the server can simultaneously handle. The default value of this flag is 16.

Use the no form of this command to set the value to the default value 16.

Command Syntax

```
max-connections <1-32>
```

Parameters

<1-32>

Set the maximum number of connections value between 1 and 32.

Default

16

Command mode

twamp-server mode

Applicability

This command was introduced in OcNOS version 6.2.0

Examples

```
OcNOS(config)#twamp server
OcNOS(config-twamp-server)#max-connections 12
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#no max-connections
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#
```

© 2026 IP Infusion Inc. Proprietary

twamp server maximum-sessions-per-connection

Use this command to set the maximum number of sessions that the server can simultaneously accept on the same connection. The default value of this flag is 1.

Use the no form of this command to set the value to the default value 1.

Command Syntax

```
maximum-sessions-per-connection <1-2000>
```

Parameters

1-2000

Set the maximum number of sessions per connection between 1 and 2000.

Notes: The maximum number of TWAMP server sessions varies by hardware platform:

- Q2C, J2C+, and Q2C+ boards: 1–2000
- Q2U and Q2A boards: 1–1000
- Q1 boards: Channels 1–64

Default

1

Command mode

twamp-server mode

Applicability

This command was introduced in OcNOS version 6.2.0

Examples

```
OcNOS(config)#twamp server
OcNOS(config-twamp-server)#maximum-sessions-per-connection 26
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#no maximum-sessions-per-connection
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#
```

© 2026 IP Infusion Inc. Proprietary

twamp client maximum-connections

Use this command to set the maximum number of connections that the client can create. The default value of this flag is 64.Use the no form of this command to set the value to the default value 64.

Command Syntax

```
max-connections (<1-64>)
```

Parameters

1-64
Set the maximum number of connections value between 1 and 64.

Default

64

Command mode

twamp-client mode

Applicability

This command was introduced in OcNOS version 6.4.1

Examples

```
OcNOS(config)#twamp client
OcNOS(config-twamp-client)# maximum-connections 10
OcNOS(config-twamp-client)#commit
OcNOS(config-twamp-client)#no max-connections
OcNOS(config-twamp-client)#commit
```

© 2026 IP Infusion Inc. Proprietary

twamp client maximum-sessions

Use this command to set the maximum number of test sessions that the client can simultaneously handle.

Use the no form of this command to set the value to the default value 64.

Command Syntax

```
maximum-sessions (1-2000)
```

Parameters

1-2000

Set the maximum number of sessions between 1 and 2000.

Notes: The maximum number of TWAMP client sessions varies by hardware platform:

- Q2C, J2C+, and Q2C+ boards: 1–2000
- Q2U and Q2A boards: 1–1000
- Q1 boards: Channels 1–64

Default

64

Command mode

twamp-client mode

Applicability

This command was introduced in OcNOS version 6.4.1

Examples

```
OcNOS(config)#twamp client
OcNOS(config-twamp-client )#maximum-sessions 44
OcNOS(config-twamp-client)#commit
OcNOS(config-twamp-client )#no maximum-sessions
OcNOS(config-twamp-client )#commit
OcNOS(config-twamp-client )#
```

© 2026 IP Infusion Inc. Proprietary

twamp client maximum-sessions-per-connection

Use this command to set the maximum number of test sessions that the client can simultaneously create on the same connection. The default value of this flag is 1.

Use the no form of this command to set the value to the default value 1.

Command Syntax

```
maximum-sessions-per-connection (1-2000)
```

Parameters

1-2000

Set the maximum number of sessions per connection between 1 and 2000.

Notes: The maximum number of TWAMP client sessions varies by hardware platform:

- Q2C, J2C+, and Q2C+ boards: 1–2000
- Q2U and Q2A boards: 1–1000
- Q1 boards: Channels 1–64

Default

1

Command mode

twamp-client mode

Applicability

This command was introduced in OcNOS version 6.4.1

Examples

```
OcNOS(config)#twamp client
OcNOS(config-twamp-client )#maximum-sessions-per-connection 26
OcNOS(config-twamp-client )#commit
OcNOS(config-twamp-client )#no maximum-sessions-per-connection
OcNOS(config-twamp-client)#commit
OcNOS(config-twamp-client)#
```

© 2026 IP Infusion Inc. Proprietary

twamp client server

Use this command to configure the server ip address(ipv4/ipv6), hostname, and port number.

Command Syntax

```
server (ipv4 A.B.C.D|ipv6 X:X::X:X|HOSTNAME) (port (862|<1025-65535>)|)
```

Parameters

ipv4 A.B.C.D
Specifies the target server's IPv4 address

ipv6 X:X::X:X
Specifies the target server's IPv6 address

HOSTNAME
Specifies the target server's hostname

port (862|<1025-65535>
Specifies the port number on which the TWAMP server is listening for control connections. 862 is the default port.

Default

None

Command mode

twamp-client connection mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
OcNOS(config)#twamp client
OcNOS(config_twamp_client)#connection con4
OcNOS(config-twamp-connection)#server ipv4 192.168.1.100 port 862
OcNOS(config-twamp-connection)#commit

OcNOS(config)#twamp client
OcNOS(config_twamp_client)#connection con6
OcNOS(config-twamp-connection)#server ipv6 2001:0db8:85a3:0000:0000:8a2e:0370:7334 port 5000
OcNOS(config-twamp-connection)#commit

OcNOS(config)#twamp client
OcNOS(config_twamp_client)#connection conhost
OcNOS(config-twamp-connection)#server pe1-host port 5100
OcNOS(config-twamp-connection)#commit
```

© 2026 IP Infusion Inc. Proprietary

twamp client session

Use this command to request a session from the server.

Command Syntax

```
session (session-name) sender-port (port) receiver-port (port) (sender-address
<address>|) (receiver-address (address)|) (start-time (TIME)|) (interface
(IFNAME)|)
```

Parameters

session-name
Specifies the name of the TWAMP client session.

port
Specifies the sender and receiver port numbers to be used in the session.

address
Specifies the sender and receiver addresses(optional).

In case the receiver-address is not specified, the configured server address used for the connection is used for this session.

In case the sender address is not specified, the client address used for the connection is used for this session.

TIME
Defines the start time of the session (optional).

IFNAME
Specifies the interface to be associated with the session (optional).

Default

None

Command mode

twamp-client-connection mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
OcNOS(config)#twamp client
OcNOS(config_twamp_client)#connection testipv6
OcNOS(config-twamp-connection)#server ipv6 2001:67c:2468::122:e:5 port 862 OcNOS(config-twamp-
connection)#session session_IPV6 sender-port 5666 receiver-port 5666 sender-address ipv6
2001:67c:2468::122:e:5 receiver-address ipv6 2001:67c:2468::122:e:4

OcNOS(config)# twamp client
OcNOS(config_twamp_client)#connection test_ipv4 vrf-name vrf1
OcNOS(config_twamp_connection)#server ipv4 24.1.1.2 port 862
OcNOS(config-twamp-connection)#session session_IPV4 sender-port 5666 receiver-port 5666 sender-
address ipv4 24.1.1.1 receiver-address ipv4 24.1.1.2
```

© 2026 IP Infusion Inc. Proprietary

twamp-light control

Use this command to initialize the TWAMP light Control or Sender Session.

Use the `no` form of this command to un-initialize and to remove the configured TWAMP light control configurations.

Command Syntax

```
twamp-light control
no twamp-light control
```

Parameters

None

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 3.0.

Example

```
#configure terminal
(config)#twamp-light control
```

twamp-light reflector

Use this command to initialize the TWAMP reflector.

Use the `no` form of this command to un-initialize and to remove the configured TWAMP light refelctor configurations.

Command Syntax

```
twamp-light reflector
no twamp-light reflector
```

Parameters

None

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 3.0.

Examples

```
#configure terminal
(config)#twamp-light reflector
```

© 2026 IP Infusion Inc. Proprietary

twamp server admin state

Use this command to Start/Stop the TWAMP server with the above configured parameters.

In case any of the twamp server parameters are modified after the server is started, the modification will be applied after the server is stopped and started again.

Command Syntax

```
server-admin-state (enable|disable)
```

Parameters

enable
Start the TWAMP server

disable
Stop the TWAMP server

Default

Disable

Command mode

twamp-server mode

Applicability

This command was introduced in OcNOS version 6.2.0

Examples

```
OcNOS(config)#twamp server
OcNOS(config-twamp-server)#server-admin-state enable
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#server-admin-state disable
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#a
```

© 2026 IP Infusion Inc. Proprietary

twamp server max-connections-per-client

Use this command to set the maximum number of connections that the server can simultaneously handle from the same client (same IP). The default value of this flag is 1.

Use the no form of this command to set the value to the default value 1.

Command Syntax

```
maximum-connections-per-client <1-4>
```

Parameters

<1-4>

Set the maximum number of connections per client value between 1 and 4.

Default

Default value is 1

Command mode

twamp-server mode

Applicability

This command was introduced in OcNOS version 6.2.0

Examples

```
OcNOS(config)#twamp server
OcNOS(config-twamp-server)#max-connections-per-client 3
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#no max-connections-per-client
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#
```

© 2026 IP Infusion Inc. Proprietary

twamp server maximum-sessions

Use this command to set the maximum number of sessions that the server can simultaneously handle. The default value of this flag is 64.

Use the no form of this command to set the value to the default value 64.

Command Syntax

```
maximum-sessions <1-2000>
```

Parameters

1-2000

Set the maximum number of sessions between 1 and 2000.

Notes: The maximum number of TWAMP server sessions varies by hardware platform:

- Q2C, J2C+, and Q2C+ boards: 1–2000
- Q2U and Q2A boards: 1–1000
- Q1 boards: Channels 1–64

Default

Default value is 64

Command mode

twamp-server mode

Applicability

This command was introduced in OcNOS version 6.2.0

Examples

```
OcNOS(config)#twamp server
OcNOS(config-twamp-server)#maximum-sessions 44
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#no max
OcNOS(config-twamp-server)#no maximum-sessions
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#
```

© 2026 IP Infusion Inc. Proprietary

twamp server port

Use this command to set the twamp server mode TCP listen port.

Use the no form of this command to reset the port value to the default value 862.

Command Syntax

```
port (862|<1025-65535>)
```

Parameters

<1025-65535>

Set the TCP port value where the twamp server will be listening. It could be either the default value 862, or in the range 1025-65535.

Default

862

Command mode

twamp-server mode

Applicability

This command was introduced in OcNOS version 6.2.0.

Examples

```
OcNOS(config)#twamp server
OcNOS(config-twamp-server)#port 2555
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#no port
OcNOS(config-twamp-server)#commit
OcNOS(config-twamp-server)#
```

© 2026 IP Infusion Inc. Proprietary

twamp start-test-session (Deprecated)

This command is not supported anymore. The “delay-measurement dynamic twamp” command under interface mode should be used to start twamp measurement instead.

Use this command to start the TWAMP test session.

Command Syntax

```
twamp start test-session WORD packet-count <1-65535> (interval <500-10000>|) (delay-threshold <0-
1000000|)
```

Parameters

WORD
TWAMP test session name.

<1-65535>
Number packed to send for test.

<500-10000>
Interval to send each packet in milliseconds.

<0-1000000>
Delay threshold, in microseconds (default = 0). A notification is sent when the average delay crosses this value during the measurement. However, no checking is carried on the delay value when it is set to 0, hence, no notification is sent.

Default

1000 for interval

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 3.0.

Examples

```
#twamp start test-session s1 packet-count 10 interval 500 delay-threshold 100
```

© 2026 IP Infusion Inc. Proprietary

test-session-name (Deprecated)

This command is not supported anymore. The “delay-measurement dynamic twamp” command under interface mode should be used to start twamp measurement instead.

Use this command to configure the TWAMP test sessions.

Use the `no` form of this command to un-configure or remove the configured TWAMP test sessions.

Notes:

- The dscp-to-queue mapping is performed for egress only when the sender port has the default value 862. Otherwise, Queue-6 will always be used.
- The dscp-to-queue mapping is done based on the first 3 bits of the `dscp` value.
- If `dscp` is not supplied (value 0), Queue-6 will always be used.

Command Syntax

```
test-session-name WORD sender-ip (A.B.C.D|X:X::X:X|HOSTNAME) (sender-port (862|<1025-65535>)|)
reflector-ip (A.B.C.D|X:X::X:X|HOSTNAME) (reflector-port (862|<1025-65535>)|)(dscp
(default|cs1|af11|af12|af13|cs2|af21|af22|af23|cs3|af31|af32|af33|cs4|af41|af42|af43|cs5|ef|cs6|cs7|<
0-63>)|)
no test-session-name WORD
```

Parameters

WORD
TWAMP test session name.

X:X::X:X
TWAMP Sender or Reflector IP address in IPv6 format.

A.B.C.D
TWAMP Sender or Reflector IP address in IPv4 format.

<1025-65535>
TWAMP Sender or Reflector UDP port number.

dscp
The DSCP value to used by the sender for TWAMP messages. This can be an explicit value 0-63, or it can be a well-known name:

`(af11| af12| af13| af21| af22| af23| af31| af32| af33| af41| af42| af43| cs1| cs2| cs3| cs4| cs5| cs6| cs7| default| ef)`

Default

The `sender-port` and `reflector-port` default values are 862.

The `dscp` default value is 0.

Command Mode

TWAMP light control mode (config-twamp-light-ctrl)

Applicability

This command was introduced before OcNOS version 3.0. The `dscp` parameter was added in OcNOS version 5.0.

© 2026 IP Infusion Inc. Proprietary

Examples

To configure:

```
(config)#twamp-light control
(config-twamp-light-ctrl)# test-session-name s1 sender-ip 10.0.0.1 sender-port 1234 reflector-ip
10.0.0.1 reflector-port 1234 dscp af23
```

To un-configure:

```
(config)#twamp-light control
(config-twamp-light-ctrl)# no test-session-name s1
```

© 2026 IP Infusion Inc. Proprietary

VRF LITE CONFIGURATION

RIP Configuration 2881
Topology 2881
Configuration IPv4 VRF 2881
Validation 2882
Configuration IPv6 VRF 2883
Validation 2885
VRF Configuration 2886
Overview 2886
Topology 2886
Default VRF 2886
User-Defined VRF 2886
Validation 2887
BGP VRF Configuration 2889
Overview 2889
Configurations 2889
Validation 2890
Static Inter-VRF Route Leaking Configuration 2893
Overview 2893
Topology 2893
IPv4 Inter-VRF Route Leaking Configuration 2893
IPv6 Inter-VRF Route Leaking Configuration 2895
OSPF VRF Configuration 2897
Overview 2897
Topology 2897
Configuration IPv4 VRF 2897
IS-IS VRF Configuration 2899
Overview 2899
Topology 2899
Configuration IPv4 VRF 2899
IS-IS IPv6 VRF Configuration 2902
Overview 2902
Configuration 2902
Maximum Allowed VRF FIB Routes in BGP/MPLS L3VPN Configuration 2905
Overview 2905
Feature Characteristics 2905
Benefits 2905
Configuration 2906

© 2026 IP Infusion Inc. Proprietary

Maximum FIB Routes Commands 2920
Glossary 2924

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 194. RIP topology for VR/VRF

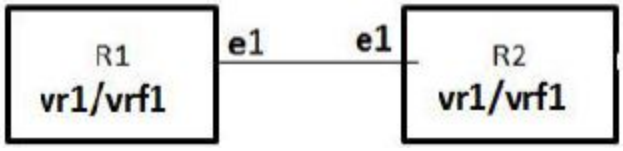

Configuration IPv4 VRF

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.
`#configure terminal`	Enter configure 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

© 2026 IP Infusion Inc. Proprietary

`(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)#exit`	Exit router mode.
`(config)#interface eth1`	Exit 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.

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.
`#configure terminal`	Enter configure 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)#exit`	Exit router mode.
`(config)#interface eth1`	Exit 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 2.2.2.2 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
```

© 2026 IP Infusion Inc. Proprietary

```
             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
```

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
```

Configuration IPv6 VRF

R1

`#configure terminal`	Enter configure mode.
`(config)#virtual-router VR1`	Create virtual router VR1.
`(config-vr)#load ipv6 rip`	Load the RIPng module to be used in VR1.
`(config)#exit`	Exit VR mode.
`(config)#interface eth1`	switch to 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.
`#configure terminal`	Enter configure mode.
`(config)#ip vrf vrf1`	Create vrf1 .

© 2026 IP Infusion Inc. Proprietary

(config-vrf)#exit	Exit VRF mode.
(config)#router ipv6 rip	Enable ipv6 rip
(config-router)#version 2	Specify RIP version 2.
(config-router)#address-family ipv4 vrf vrf1	Enter address family mode for vrf1
(config-router-af)#aggregate-address 2222::/48	Configure the ipv6 aggregate-address
(config-router-af)#exit	Exit address-family mode.
(config-router)#exit	Exit router mode.
(config)#interface eth1	Exit 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)#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 configure mode.
(config)#virtual-router VR1	Create virtual router VR1.
(config-vr)#load ipv6 rip	Load the RIPng module to be used in VR1.
(config)#exit	Exit VR mode.
(config)#interface eth1	switch to 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.
#configure terminal	Enter configure mode.
(config)#ip vrf vrf1	Create vrf1 .
(config-vrf)#exit	Exit VRF mode.
(config)#router ipv6 rip	Enable ipv6 rip
(config-router)#version 2	Specify RIP version 2.
(config-router)#address-family ipv4 vrf vrf1	Enter address family mode for vrf1
(config-router-af)#aggregate-address 2222::/48	Configure the ipv6 aggregate-address
(config-router-af)#exit	Exit address-family mode.
(config-router)#exit	Exit router mode.
(config)#interface eth1	Exit interface mode.
(config-if)ip vrf forwarding vrf1	Associate the interface to vrf1.

© 2026 IP Infusion Inc. Proprietary

`(config-if)#ip address 2.2.2.2/24`	Configure the IP address 2.2.2.2 to 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

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
```

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 195. 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)#commit`	Commit the candidate configuration to the running configuration
`(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
`(config-if)#commit`	Commit the candidate configuration to the running configuration

User-Defined VRF

`#con t`	Enter configuration mode.
`(config)#ip vrf vrf1`	Create vrf1. Use no ip vrf vrf1 to delete the created vrf

© 2026 IP Infusion Inc. Proprietary

(config)#exit	Exit configure mode.
(config)#interface eth1	Enter interface mode
(config-if)#ip vrf forwarding vrf1	Associate eth1 to vrf1. Use no ip vrf forwarding vrf1 for un-configuration
(config-if)#ip address 3.3.3.2/24	Configure the IP address 3.3.3.2 to eth1
(config-if)#commit	Commit the candidate configuration to the running configuration
(config-if)#exit	Exit interface mode.

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
(config-if)#commit	Commit the candidate configuration to the running configuration

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
```

To display the IP routing table associated with a VRF, use the show ip route vrf vrf-name command

```
#sh 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            3.3.3.0/24 is directly connected, eth1, 00:01:22
      S            20.20.20.0/24 [1/0] is directly connected, eth1, 00:00:13

      Gateway of last resort is not set
```

Use this show command to display the static routes configured.

```
#sh ip route vrf  vrf1 static
     IP Route Table for VRF "vrf1"
     S            20.20.20.0/24 [1/0] is directly connected, eth1, 00:01:13

     Gateway of last resort is not set
```

BGP VRF 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: Enabling dynamic capability on the specific BGP peer will prevent BGP session to reset.

Configurations

Note: Users can configure up to 128 export route-targets under "IP VRF mode”.

Topology

Figure 196. BGP topology for VRF

R1

`#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 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

© 2026 IP Infusion Inc. Proprietary

(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

R1

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: ::
```

© 2026 IP Infusion Inc. Proprietary

```
BGP connection: non shared network
Last Reset: 00:00:12, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)
```

```
#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

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
Last Reset: 00:02:41, due to Configuration Change (Cease Notification sent)
Notification Error Message: (Cease/Other Configuration Change.)
```

```
#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
```

Verify BGP Adjacency

Use the following show commands to verify the BGP adjacency:

```
show ip bgp neighbor
```

© 2026 IP Infusion Inc. Proprietary

Static Inter-VRF Route Leaking Configuration

Overview

Inter-VRF route leaking allows the leaking of routes from one VRF (the source VRF) to another VRF (the destination VRF) on the same router. Inter-VRF routes can exist in any VRF (including the default and management VRF) on the system. Route-Leaking between two non-default VRF is achieved using MP-BGP (RT and RD). This section covers leaking between the Default VRF and any Non-default VRF using Statically.

Topology

Note: By default, eth0 is in management VRF.

Figure 197. Device topology

IPv4 Inter-VRF Route Leaking Configuration

R1 Configuration

`#configure terminal`	Enter the router configuration mode
`(config)#ip vrf vrf1`	Create vrf1
`(config-vrf)#exit`	Exit VRF mode
`(config)#interface eth1`	Switch to interface eth1
`(config-if)#ip address 20.1.1.1/24`	Assign the IP address 20.1.1.1 to eth1
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Switch to interface eth2.
`(config-if)#ip vrf forwarding vrf1`	Associate eth1 to vrf1. Use `no ip vrf forwarding vrf1` for un-configuration.
`(config-if)#ip address 10.1.1.1/24`	Assign the IP address 10.1.1.1 to eth2 in vrf1.
`(config-if)#exit`	Exit interface mode.
`(config)#ip route 10.1.1.0/24 eth2`	Leaking non-default VRF's connected route into default VRF.
`(config)#ip route vrf management 30.1.1.0/24 20.1.1.2 eth1 global`	Leaking a route from the Default VRF into the "management" VRF statically.
`(config)#ip route vrf vrf1 20.1.1.0/24 20.1.1.2 eth1`	Leaking a route from the Default VRF into VRF

© 2026 IP Infusion Inc. Proprietary

`global`	"vrf1" statically.
`(config)#ip route vrf vrf1 30.1.1.0/24 20.1.1.2 eth1 global`	Leaking a route from the Default VRF into VRF "vrf1" statically.
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode

R2 Configuration

`#configure terminal`	Enter the router configuration mode
`(config)#interface eth1`	Switch to interface eth1
`(config-if)#ip address 20.1.1.2/24`	Assign the IP address 20.1.1.2 to eth1 in vrf1
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Switch to interface eth2
`(config-if)#ip address 30.1.1.1/24`	Assign the IP address 30.1.1.1 to eth2 in vrf1
`(config-if)#exit`	Exit interface mode.
`(config)# ip route 10.1.1.0/24 20.1.1.1`	Configuring static route to non-default VRF.
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode

Validation

R1

```
#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"
S    v       10.1.1.0/24 [1/0] is directly connected, xe12, 00:37:31
C            20.1.1.0/24 is directly connected, xe14, 00:00:17
C            127.0.0.0/8 is directly connected, lo, 3d06h01m
IP Route Table for VRF "management"
C            10.12.87.0/24 is directly connected, eth0, 3d06h00m
S    v       30.1.1.0/24 [1/0] via 20.1.1.2, xe14, 00:00:17
C            127.0.0.0/8 is directly connected, lo.management, 3d06h01m
IP Route Table for VRF "vrf1"
C            10.1.1.0/24 is directly connected, xe12, 02:27:07
S    v       20.1.1.0/24 [1/0] via 20.1.1.2, xe14, 00:00:17
S    v       30.1.1.0/24 [1/0] via 20.1.1.2, xe14, 00:00:17
C            127.0.0.0/8 is directly connected, lo.vrf1, 02:41:58
```

To display the IP routing table associated with a VRF, use the `show ip route vrf vrf-name` command

```
#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, E - EVPN,
       v - vrf leaked
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
S    *> v10.1.1.0/24 [1/0] is directly connected, xe12, 00:39:15
C    *>  20.1.1.0/24 is directly connected, xe14, 00:02:01
C    *>  127.0.0.0/8 is directly connected, lo, 3d06h02m
IP Route Table for VRF "management"
C    *>  10.12.87.0/24 is directly connected, eth0, 3d06h02m
S    *> v30.1.1.0/24 [1/0] via 20.1.1.2, xe14, 00:02:01
C    *>  127.0.0.0/8 is directly connected, lo.management, 3d06h02m
IP Route Table for VRF "vrf1"
C    *>  10.1.1.0/24 is directly connected, xe12, 02:28:51
S    *> v20.1.1.0/24 [1/0] via 20.1.1.2, xe14, 00:02:01
S    *> v30.1.1.0/24 [1/0] via 20.1.1.2, xe14, 00:02:01
C    *>  127.0.0.0/8 is directly connected, lo.vrf1, 02:43:42

Gateway of last resort is not set
```

IPv6 Inter-VRF Route Leaking Configuration

IPv6 Statically Routes can be leaked between the Default VRF and Private VRF and vice-versa.

Note: IPv6 Leaking between "management" VRF and the default VRF is not supported.

R1 Configuration

`#con t`	Enter the router configuration mode
`(config)#ip vrf vrf1`	Create vrf1
`(config-vrf)#exit`	Exit VRF mode
`(config)#interface eth1`	Switch to interface eth1
`(config-if)#ipv6 address 2001::1/64`	Assign the IPv6 address 2001::1 to eth1
`(config-if)#exit`	Exit interface mode.
`(config)#interface eth2`	Switch to interface eth2.
`(config-if)#ip vrf forwarding vrf1`	Associate eth1 to vrf1. Use no ip vrf forwarding vrf1 for un-configuration.
`(config-if)#ipv6 address 3001::1/64`	Assign the IPv6 address 3001::1 to eth2 in vrf1
`(config-if)#exit`	Exit interface mode.
`(config)# ipv6 route 3001::/64 eth2`	Leaking non-default VRF's connected route into default VRF.

© 2026 IP Infusion Inc. Proprietary

`OcNOS(config)#ipv6 route 3333::2/128 fe80::da9e:f3ff:fea2:6ce6 eth2`	Leaking non-default VRF's non-connected route into default VRF
`OcNOS(config)#ipv6 route vrf vrf1 2001::/64 eth1 global`	Leaking a connected route from the Default VRF into VRF "vrf1" statically.
`OcNOS(config)#ipv6 route vrf vrf1 2222::1/128 fe80::da9e:f3ff:fec9:65a4 eth1 global`	Leaking a route from the Default VRF into the "vrf1" VRF statically
`(config-if)#commit`	Commit the candidate configuration to the running configuration
`(config-if)#exit`	Exit interface mode

Validation

R1

```
OcNOS#sh  run ipv6 route
!
ipv6 route 3001::/64 eth2
ipv6 route 3333::/128 fe80::da9e:f3ff:fea2:6ce6 eth2
ipv6 route vrf vrf1 2001::/64 eth1 global
ipv6 route vrf vrf1 2222::1/128 fe80::da9e:f3ff:fec9:65a4 eth1 global
!
OcNOS#

OcNOS#sh 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
      v - vrf leaked
Timers: Uptime


IP Route Table for VRF "default"
C       ::1/128 via ::, lo, 01:14:24
C       2001::/64 via ::, eth1, 01:14:24
O       2222::1/128 [110/1] via fe80::da9e:f3ff:fec9:65a4, eth1, 01:13:30
S   v   3001::/64 [1/0] via ::, eth2, 00:00:32
S   v   3333::2/128 [1/0] via fe80::da9e:f3ff:fea2:6ce6, eth2, 00:01:22
C       fe80::/64 via ::, eth4, 01:14:24
IP Route Table for VRF "management"
C       ::1/128 via ::, lo.management, 01:14:24
C       fe80::/64 via ::, eth0, 01:14:24
IP Route Table for VRF "vrf1"
C       ::1/128 via ::, lo.vrf1, 01:14:24
S   v   2001::/64 [1/0] via ::, eth1, 01:14:24
S   v   2222::1/128 [1/0] via fe80::da9e:f3ff:fec9:65a4, eth1, 01:14:24
C       3001::/64 via ::, eth2, 01:14:24
i L1    3333::2/128 [115/20] via fe80::da9e:f3ff:fea2:6ce6, eth2, 01:14:17
C       fe80::/64 via ::, eth2, 01:14:24
```

© 2026 IP Infusion Inc. Proprietary

OSPF VRF 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 198. 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
`(config-if)#commit`	Commit the candidate configuration to the running configuration

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`	Specify the network type and area 0.

© 2026 IP Infusion Inc. Proprietary

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
(config-if)#commit	Commit the candidate configuration to the running configuration

Validation

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, 00:02:27

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, 00:01:49

Gateway of last resort is not set
```

IS-IS VRF 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 199. ISIS Topology for VRF

Configuration IPv4 VRF

R1

`#configure terminal`	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)#commit`	Commit the candidate configuration to the running configuration

© 2026 IP Infusion Inc. Proprietary

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

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 : vrf1
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        5254.0047.dc01      Up     24        L1   IS-IS
```

R2

```
#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 0
Total number of adjacencies: 1
Tag 1:  VRF : vrf1
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0001 eth1        5254.0060.d535      Up     5         L1   IS-IS
```

© 2026 IP Infusion Inc. Proprietary

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 : vrf1
     Destination         Metric      Next-Hop            Interface      Tag
C    2.2.2.0/24          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 : vrf1
     Destination         Metric      Next-Hop            Interface      Tag
C    2.2.2.0/24          10          -                  eth1            0
```

© 2026 IP Infusion Inc. Proprietary

IS-IS IPv6 VRF 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.

Configuration

Topology

Figure 200. ISISv6 Topology for VRF

R1

`#configure terminal`	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)#ipv6 address 1000::1/64`	Configure ipv6 address on eth1
`(config-if)#ipv6 router isis 1`	Enable ISISv6 routing on an instance for area 49
`(config-if)#isis circuit-type level-1`	Configure interface as level-1
`(config-if)#commit`	Commit the candidate configuration to the running configuration

© 2026 IP Infusion Inc. Proprietary

(config-if)#exit	Exit interface mode
(config)#exit	Exit config mode

R2

#configure terminal	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)#ipv6 address 1000::2/64	Configure ipv6 address on eth1
(config-if)#ipv6 router isis 1	Enable ISISv6 routing on an instance for area 49
(config-if)#isis circuit-type level-1	Configure interface as level-1
(config-if)#commit	Commit the candidate configuration to the running configuration
(config-if)#exit	Exit interface mode
(config)#exit	Exit config mode

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 : vrf1
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0000.0002 eth1        b86a.97c4.31c5      Up     27        L1   IS-IS
R1#
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 1:  VRF : vrf1
C    1000::/64 [10]
      via ::, eth1

R1#

#
R2#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 0
Total number of adjacencies: 1
```

```
Tag 1:  VRF : vrf1
System Id      Interface   SNPA                State  Holdtime  Type Protocol
0000.0000.0001 eth1        b86a.97cb.3ec5      Up     8         L1   IS-IS
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 1:  VRF : vrf1
C    1000::/64 [10]
      via ::, eth1

R2#
```

Maximum Allowed VRF FIB Routes in BGP/MPLS L3VPN Configuration

Overview

Virtual Routing and Forwarding (VRF) in BGP/MPLS L3VPN enables multiple isolated routing tables within a single physical router, ensuring secure and independent routing for different customers. VRF instances maintain routing and forwarding information by importing routes from connected Customer Edge (CE) devices and remote VPN sites via MP-BGP PE routers. Each VRF can run its own routing protocols, such as OSPF, IS-IS, or E-BGP, enhancing isolation.

To optimize routing resources, a maximum route limit can be configured per VRF. This provides ability to control the number of imported dynamic routes that appear in the forwarding information base (FIB). When this limit is reached, new dynamic routes remain pending, while static and connected routes continue contributing to the total count. A threshold-based NetConf notification and syslog warning message can notify administrators when the route count approaches the limit, ensuring efficient resource management and maintaining network stability.

Feature Characteristics

This feature enhances BGP/MPLS L3VPN scalability by allowing per-VRF route limits, ensuring independent routing tables for each VPN instance. It applies limits only to dynamic routes while accounting for static and connected routes in the total count. A threshold-based NetConf notification and syslog warning message notifies administrators when the route count nears the limit and clears once it drops below, ensuring efficient resource management.

- Per-VRF Route Limit Configuration: Allows administrators to set a maximum number of dynamic routes that can be installed in the VRF FIB.
- Isolation of Routing Tables: Ensures that each VRF instance maintains its own independent routing table without interference from other VPN instances.
- Support for Dynamic, Static, and Connected Routes: The maximum route limit applies only to dynamic routes, while static and connected routes contribute to the total route count.
- Threshold-Based NetConf Notification and Syslog Warning Message: Generates a NetConf notification and syslog warning message when the number of installed routes reaches a configurable percentage of the maximum limit. The NetConf notification and syslog warning message clears once the count drops below the threshold.

Benefits

- Optimized Resource Utilization: Prevents a single VRF from consuming excessive hardware routing table space, ensuring fair allocation across multiple VPN instances.
- Enhanced Network Stability: Enforces a configurable route limit to prevent routing table exhaustion and unexpected failures.
- Improved Scalability: Supports large-scale multi-tenant deployments by efficiently managing routing resources.

© 2026 IP Infusion Inc. Proprietary

- Proactive Monitoring: Provides configurable threshold alerts, allowing administrators to take preventive actions before reaching critical limits.
- Better Control Over Dynamic Routes: Ensures only necessary routes are installed in the VRF FIB while maintaining seamless routing protocol operations.
- Reliable and Predictable Network Performance: Helps maintain a well-balanced network infrastructure with controlled resource usage.
- Streaming telemetry support enables continuous monitoring of the current utilization percentage and threshold exceedance status for the VRF FIB route limit, allowing administrators to take proactive actions.

Configuration

Configure the maximum number of VRF FIB routes on a networking device.

Topology

This topology represents a basic MPLS Layer 3 VPN (L3VPN) setup using VRF instances and IS-IS as the Interior Gateway Protocol (IGP). The topology consists of two Customer Edge (CE) routers, CE1 and CE2, connected to a single Provider Edge (PE1) router via different virtual routing instances. CE1 is connected to PE1 through interface xe13.3, which is associated with VRF3, while CE2 connects via xe1.3, mapped to VRF3.

Figure 201. Max VRF FIB Routes Topology

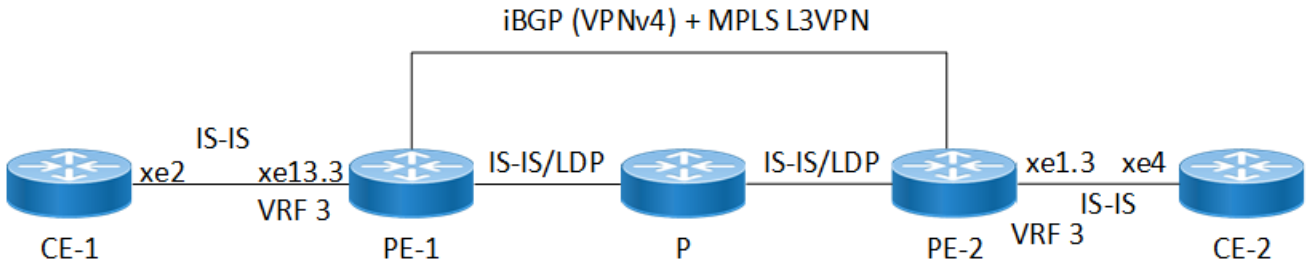

The following configuration on the PE1 router defines a VRF instance (VRF3) with a maximum limit of 10 dynamic IPv4 FIB routes, applying a stop-install policy when the limit is reached and setting an 70% threshold for early warnings. It also configures interface xe13.3 with VLAN encapsulation and an IP address, ensuring proper traffic segmentation and routing within the VRF. Follow the steps below to configure the maximum number of VRF FIB routes on a networking device.

Notes:

- The limit applies only to dynamic routes and does not apply to static or connected routes.
- If recursive routes exist and the configured limit is lower than the current route count in the FIB, the resolved route information for recursive routes may be incomplete in the "`show ip route`" command output.
- Execute `clear ip route <vrf-name>` or `clear ipv6 route <vrf-name>` when the FIB route count already exceeds the configured limit.

Configure PE1 router as follows:

1. To enter into config mode, execute the following command in the configure mode.

```
PE1# configure terminal
```

2. Create a VRF.

```
PE1 (config)#ip vrf VRF3
PE1 (config-vrf)#rd 11.11.11.3:3
PE1 (config-vrf)#route-target both 100:3
```

3. Set the maximum limit of 10 IPv4 FIB routes to stop installing new routes after reaching the limit.

```
 PE1 (config-vrf) #maximum-fib-routes ipv4 10 stop-install
```

4. It triggers a warning log once the limit is reached.

```
PE1 (config-vrf) #maximum-fib-routes ipv4 10 warning-only
```

5. Set a threshold of 70% for IPv4 FIB routes.

```
PE1 (config-vrf) #maximum-fib-routes ipv4 threshold 70
```

6. Set the maximum limit of 10 IPv6 FIB routes to stop installing new routes after reaching the limit.

```
PE1 (config-vrf) #maximum-fib-routes ipv6 10 stop-install
```

7. It triggers a warning log once the limit is reached.

```
 PE1 (config-vrf) #maximum-fib-routes ipv6 10 warning-only
```

8. Set a threshold of 70% for IPv6 FIB routes.

```
PE1 (config-vrf) #maximum-fib-routes ipv6 threshold 70
```

9. Commit the changes and exit.

```
 PE1 (config-vrf)#commit
 PE1 (config-vrf)#exit
```

10 Enter the interface xe13.3 to be configured and assign the Interface to a VRF.

```
PE1(config)#interface xe13.3
PE1(config-if)#ip vrf forwarding VRF3
```

11 Enable VLAN Encapsulation (Dot1Q).

```
PE1(config-if)#encapsulation dot1q 3
```

12 Configure an IP address to the interface.

```
PE1(config-if)#i p address 10.1.13.1/24
```

13 Commit the changes and exit.

```
PE1(config-if)# commit
PE1(config-if)# exit
```

© 2026 IP Infusion Inc. Proprietary

Running Configuration on PE1 Router

The running configuration on PE1 sets VRF3 with a 10-route FIB limit, an 80% threshold warning, and assigns VLAN encapsulation and an IP to interface xe13.3.

```
ip vrf VRF3
 rd 11.11.11.3:3
 route-target both 100:3
 maximum-fib-routes ipv4 10 stop-install
 maximum-fib-routes ipv4 threshold 70
 maximum-fib-routes ipv6 10 stop-install
 maximum-fib-routes ipv6 threshold 70
!
# Sub interface config
interface xe13.3
 description L3VPN-6VPE-VRF3
 encapsulation dot1q 3
 load-interval 30
 ip vrf forwarding VRF3
 ip address 10.1.13.1/24
 ipv6 address 2001:0:3::1/64
 isis network point-to-point
 ip router isis CEPE3
 ipv6 router isis CEPE3
!
# ISIS config
router isis CEPE3 VRF3
 is-type level-1-2
 metric-style wide
 net 49.0001.0000.0000.0003.00
 redistribute bgp
 !
 address-family ipv6
 redistribute bgp
 exit-address-family
!
# BGP config for VRF
router bgp 1001
 bgp router-id 12.12.12.1
 bgp auto-policy-soft-reset enable
 bgp graceful-restart
 neighbor 14.14.14.1 remote-as 1001
 neighbor 14.14.14.1 update-source lo
 !
 address-family ipv4 unicast
 neighbor 14.14.14.1 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 14.14.14.1 activate
 exit-address-family
 !
 address-family ipv4 vrf VRF3
 redistribute connected
 redistribute isis
 exit-address-family
 !
 address-family ipv6 vrf VRF3
 redistribute connected
 redistribute isis
 exit-address-family
 !
```

© 2026 IP Infusion Inc. Proprietary

Validation for PE1

Validation ensures that the Max VRF FIB Routes Implementation functions as expected by verifying route limits, isolation, and NetConf notification and syslog warning message, preventing resource exhaustion and ensuring efficient routing operations.

To Check BGP VPNv4 Peering between PE1--PE2

```
PE1#show ip bgp vpnv4 all summary
BGP router identifier 12.12.12.1, local AS number 1001
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   Desc
14.14.14.1         4  1001        177        282       2      0      0  01:07:05            11

Total number of neighbors 1

Total number of Established sessions 1


#  Verify VRF Route Threshold (e.g., Max Route Limit)

PE1#show ip route vrf VRF3 summary

------------------------------------------------
IP routing table name is VRF3(5)
------------------------------------------------
IP routing table maximum-paths   : 8
Total number of IPv4 routes      : 10
Total number of IPv4 paths       : 10
Pending routes (due to route max reached): 6
Route Source    Networks
connected       4
static          2
isis            4
Total           10
FIB             10

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

VRF FIB Route Limits:
 Configured Route Limit    : 10
 Utilization Percentage    : 100 %
 Action upon reaching limit: stop-install
 Warning Threshold         : 70 %
 Exceeds Threshold         : Yes

PE1#show ip route vrf VRF3
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 "VRF3"
S            10.1.5.0/24 [1/0] is directly connected, xe13.3, installed 00:44:19, last update
01:17:46 ago
S            10.1.6.0/24 [1/0] is directly connected, xe13.3, installed 00:44:19, last update
01:17:46 ago
```

© 2026 IP Infusion Inc. Proprietary

```
C            10.1.13.0/24 is directly connected, xe13.3, installed 00:44:19, last update 01:17:45 ago
C            10.3.1.1/32 is directly connected, loopback1, installed 00:44:19, last update 01:17:45
ago
C            10.10.2.2/32 is directly connected, loopback2, installed 00:44:19, last update 01:17:45
ago
C            127.0.0.0/8 is directly connected, lo.VRF3, installed 00:44:19, last update 01:17:56 ago
i L2         192.0.1.0/24 [115/11] via 10.1.13.111, xe13.3, installed 00:44:19, last update 01:17:32
ago
i L2         192.0.2.0/24 [115/11] via 10.1.13.111, xe13.3, installed 00:44:19, last update 01:17:32
ago
i L2         192.0.3.0/24 [115/11] via 10.1.13.111, xe13.3, installed 00:44:19, last update 01:17:32
ago
i L2         192.0.4.0/24 [115/11] via 10.1.13.111, xe13.3, installed 00:44:19, last update 01:17:32
ago

Gateway of last resort is not set

PE1#sh ip route vrf VRF3 database
all        bgp        connected  isis       kernel     ospf       rip        static
PE1#show ip route vrf VRF3 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 "VRF3"
B     >  1.1.1.0/24 [200/11] via 14.14.14.1, installed 00:36:18, last update 00:36:18 ago
B     >  1.1.2.0/24 [200/11] via 14.14.14.1, installed 00:36:18, last update 00:36:18 ago
B     >  1.1.3.0/24 [200/11] via 14.14.14.1, installed 00:36:18, last update 00:36:18 ago
B     >  1.1.4.0/24 [200/11] via 14.14.14.1, installed 00:36:18, last update 00:36:18 ago
B     >  1.1.5.0/24 [200/11] via 14.14.14.1, installed 00:36:18, last update 00:36:18 ago
B     >  1.1.6.0/24 [200/11] via 14.14.14.1, installed 00:36:18, last update 00:36:18 ago
B     >  1.1.7.0/24 [200/11] via 14.14.14.1, installed 00:36:18, last update 00:36:18 ago
B     >  1.1.8.0/24 [200/11] via 14.14.14.1, installed 00:36:18, last update 00:36:18 ago
B     >  1.1.9.0/24 [200/11] via 14.14.14.1, installed 00:36:18, last update 00:36:18 ago
B     >  1.1.10.0/24 [200/11] via 14.14.14.1, installed 00:36:18, last update 00:36:18 ago
S    *>  10.1.5.0/24 [1/0] is directly connected, xe13.3, installed 00:44:29, last update 01:17:56
ago
S    *>  10.1.6.0/24 [1/0] is directly connected, xe13.3, installed 00:44:29, last update 01:17:56
ago
C    *>  10.1.13.0/24 is directly connected, xe13.3, installed 00:44:29, last update 01:17:55 ago
i L1     10.1.13.0/24 [115/10] is directly connected, xe13.3, installed 01:17:54, last update
01:17:54 ago
B     >  10.1.23.0/24 [200/0] via 14.14.14.1, installed 00:36:18, last update 00:36:18 ago
C    *>  10.3.1.1/32 is directly connected, loopback1, installed 00:44:29, last update 01:17:55 ago
C    *>  10.10.2.2/32 is directly connected, loopback2, installed 00:44:29, last update 01:17:55 ago
C    *>  127.0.0.0/8 is directly connected, lo.VRF3, installed 00:44:29, last update 01:18:06 ago
i L2 *>  192.0.1.0/24 [115/11] via 10.1.13.111, xe13.3, installed 00:44:29, last update 01:17:42 ago
i L2 *>  192.0.2.0/24 [115/11] via 10.1.13.111, xe13.3, installed 00:44:29, last update 01:17:42 ago
i L2 *>  192.0.3.0/24 [115/11] via 10.1.13.111, xe13.3, installed 00:44:29, last update 01:17:42 ago
i L2 *>  192.0.4.0/24 [115/11] via 10.1.13.111, xe13.3, installed 00:44:29, last update 01:17:42 ago
i L2  >  192.0.5.0/24 [115/11] via 10.1.13.111, xe13.3, installed 00:44:29, last update 01:17:42 ago
i L2  >  192.0.6.0/24 [115/11] via 10.1.13.111, xe13.3, installed 00:44:29, last update 01:17:42 ago
i L2  >  192.0.7.0/24 [115/11] via 10.1.13.111, xe13.3, installed 00:44:29, last update 01:17:42 ago
i L2  >  192.0.8.0/24 [115/11] via 10.1.13.111, xe13.3, installed 00:44:29, last update 01:17:42 ago
i L2  >  192.0.9.0/24 [115/11] via 10.1.13.111, xe13.3, installed 00:44:29, last update 01:17:42 ago
i L2  >  192.0.10.0/24 [115/11] via 10.1.13.111, xe13.3, installed 00:44:29, last update 01:17:42 ago

Total number of IPv4 routes 28

Gateway of last resort is not set

PE1#show rib pending-txlist ipv4 vrf VRF3
IPv4 pending:
0x02634680: 192.0.5.0/24       req(1) new_req(0) fib_notify(0) seq(0) lock(3)
```

© 2026 IP Infusion Inc. Proprietary

```
0x02634830: 192.0.6.0/24       req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x026349e0: 192.0.7.0/24       req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x02634bb0: 192.0.8.0/24       req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x02652490: 192.0.9.0/24       req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x02652660: 192.0.10.0/24      req(1) new_req(0) fib_notify(0) seq(0) lock(3)

PE1#show ip bgp vpnv4 vrf VRF3
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
Route Distinguisher: 11.11.11.3:3 (Default for VRF VRF3)
*>i  1.1.1.0/24       14.14.14.1           11       100        0    ?        -
*>i  1.1.2.0/24       14.14.14.1           11       100        0    ?        -
*>i  1.1.3.0/24       14.14.14.1           11       100        0    ?        -
*>i  1.1.4.0/24       14.14.14.1           11       100        0    ?        -
*>i  1.1.5.0/24       14.14.14.1           11       100        0    ?        -
*>i  1.1.6.0/24       14.14.14.1           11       100        0    ?        -
*>i  1.1.7.0/24       14.14.14.1           11       100        0    ?        -
*>i  1.1.8.0/24       14.14.14.1           11       100        0    ?        -
*>i  1.1.9.0/24       14.14.14.1           11       100        0    ?        -
*>i  1.1.10.0/24      14.14.14.1           11       100        0    ?        -
*> l 10.1.13.0/24     0.0.0.0              0        100        32768  ?        -
*                     0.0.0.0              10       100        32768  ?        -
*>i  10.1.23.0/24     14.14.14.1           0        100        0    ?        -
*> l 10.3.1.1/32      0.0.0.0              0        100        32768  ?        -
*> l 10.10.2.2/32     0.0.0.0              0        100        32768  ?        -
*> l 192.0.1.0        10.1.13.111          11       100        32768  ?        -
*> l 192.0.2.0        10.1.13.111          11       100        32768  ?        -
*> l 192.0.3.0        10.1.13.111          11       100        32768  ?        -
*> l 192.0.4.0        10.1.13.111          11       100        32768  ?        -
*> l 192.0.5.0        10.1.13.111          11       100        32768  ?        -
*> l 192.0.6.0        10.1.13.111          11       100        32768  ?        -
*> l 192.0.7.0        10.1.13.111          11       100        32768  ?        -
*> l 192.0.8.0        10.1.13.111          11       100        32768  ?        -
*> l 192.0.9.0        10.1.13.111          11       100        32768  ?        -
*> l 192.0.10.0       10.1.13.111          11       100        32768  ?        -
 Announced routes count = 14
 Accepted routes count = 11

PE1#ping 14.14.14.1
Press CTRL+C to exit
PING 14.14.14.1 (14.14.14.1) 100(128) bytes of data.
108 bytes from 14.14.14.1: icmp_seq=1 ttl=64 time=0.471 ms
108 bytes from 14.14.14.1: icmp_seq=2 ttl=64 time=0.351 ms
108 bytes from 14.14.14.1: icmp_seq=3 ttl=64 time=0.403 ms
108 bytes from 14.14.14.1: icmp_seq=4 ttl=64 time=0.509 ms

--- 14.14.14.1 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3087ms
rtt min/avg/max/mdev = 0.351/0.433/0.509/0.060 ms

# IPv6
PE1#show ipv6 route vrf VRF3 summary

-----------------------------------------------
IPv6 routing table name is VRF3(5)
-----------------------------------------------
IPv6 routing table maximum-paths : 8
Total number of IPv6 routes      : 10
Total number of IPv6 paths       : 10
Pending routes (due to route max reached): 4
Route Source    Networks
connected       4
isis            6
Total           10
```

```
FIB             10

ECMP statistics (active in ASIC):
 Total number of IPv6 ECMP routes   : 0
 Total number of IPv6 ECMP paths    : 0

VRF FIB Route Limits:
 Configured Route Limit    : 10
 Utilization Percentage    : 100 %
 Action upon reaching limit: stop-install
 Warning Threshold         : 70 %
 Exceeds Threshold         : Yes

PE1#show rib pending-txlist ipv6 vrf VRF3
IPv6 pending:
0x0268c600: 2000::6:0:0:0:0/64 req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x0268c7d0: 2000::7:0:0:0:0/64 req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x0268c9a0: 2000::8:0:0:0:0/64 req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x0268cb70: 2000::9:0:0:0:0/64 req(1) new_req(0) fib_notify(0) seq(0) lock(3)

# Ping VPNV4 prefixes

PE1#ping mpls l3vpn VRF3 1.1.4.0/24
Sending 5 MPLS Echos to 1.1.4.0, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

 Type 'Ctrl+C' to abort

 !
Success Rate is 100.00 percent (5/5)
PE1#ping mpls l3vpn VRF3 1.1.6.0/24
Sending 5 MPLS Echos to 1.1.6.0, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

 Type 'Ctrl+C' to abort

 !
Success Rate is 100.00 percent (5/5)
```

Running Configuration on P Router

```
# P Core Node
hostname 7004-P1
router ldp
 router-id 13.13.13.1
 transport-address ipv4 13.13.13.1
!
interface lo
 ip address 127.0.0.1/8
 ip address 13.13.13.1/32 secondary
 ipv6 address ::1/128
```

© 2026 IP Infusion Inc. Proprietary

```
 ip router isis ISIS-IGP-200
!
interface xe13
 load-interval 30
 ip address 41.41.1.2/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis ISIS-IGP-200
 enable-ldp ipv4
 mpls ldp-igp sync-delay 30
!
interface xe17
 load-interval 30
 ip address 31.31.1.2/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis ISIS-IGP-200
 enable-ldp ipv4
 mpls ldp-igp sync-delay 30
!
 exit
!
router isis ISIS-IGP-200
 is-type level-1
 net 49.0001.0100.0000.2004.00
!
!
end
```

Validation for P

Validation ensures that the Max VRF FIB Routes Implementation functions as expected by verifying route limits, isolation, and NetConf notification and syslog warning message, preventing resource exhaustion and ensuring efficient routing operations.

```
P#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
       12.12.12.1         xe13       Active    OPERATIONAL   30    01:27:49
       14.14.14.1         xe17       Passive   OPERATIONAL   30    01:18:34
P#show clns neighbors

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
Tag ISIS-IGP-200:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0100.0000.2001 xe13        b86a.97f5.eac8      Up     19        L1   IS-IS
0100.0000.2005 xe17        5c07.5851.1353      Up     21        L1   IS-IS
```

Configure PE2 router as follows:

1. To enter into config mode, execute the following command in the configure mode.

   ```
   PE2# configure terminal
   ```

2. Create a VRF.

```
PE2 (config)#ip vrf VRF3
PE2 (config-vrf)#rd 11.11.11.3:3
PE2 (config-vrf)#route-target both 100:3
```

3. Set the maximum limit of 10 IPv4 FIB routes to stop installing new routes after reaching the limit.

```
PE2 (config-vrf) #maximum-fib-routes ipv4 10 stop-install
```

4. It triggers a warning log once the limit is reached.

```
PE2 (config-vrf) #maximum-fib-routes ipv4 10 warning-only
```

5. Set a threshold of 70% for IPv4 FIB routes.

```
PE2 (config-vrf) #maximum-fib-routes ipv4 threshold 70
```

6. Set the maximum limit of 10 IPv6 FIB routes to stop installing new routes after reaching the limit.

```
PE2 (config-vrf) #maximum-fib-routes ipv6 10 stop-install
```

7. Set a threshold of 70% for IPv6 FIB routes.

```
PE2 (config-vrf) #maximum-fib-routes ipv6 threshold 70
```

8. It triggers a warning log once the limit is reached.

```
PE2 (config-vrf) #maximum-fib-routes ipv6 10 warning-only
```

9. Commit the changes and exit.

```
PE2( config-vrf)#commit
PE2 (config-vrf)#exit
```

10 Enter the interface xe1.3 to be configured and assign the Interface to a VRF.

```
PE2(config)#interface xe1.3
PE2(config-if)#ip vrf forwarding VRF3
```

11 Enable VLAN Encapsulation (Dot1Q).

```
PE2(config-if)#encapsulation dot1q 3
```

12 Configure an IP address to the interface.

```
PE2(config-if)#i p address 10.1.13.1/24
```

13 Commit the changes and exit.

```
PE2(config-if)# commit
PE2(config-if)# exit
```

Running Configuration on PE2 Router

The running configuration on PE2 sets VRF3 with a 10-route FIB limit, an 70% threshold warning, and assigns VLAN encapsulation and an IP to interface xe13.3.

© 2026 IP Infusion Inc. Proprietary

```
ip vrf VRF3
 rd 11.11.11.3:3
 route-target both 100:3
 maximum-fib-routes ipv4 10 stop-install
 maximum-fib-routes ipv4 threshold 70
 maximum-fib-routes ipv6 10 stop-install
 maximum-fib-routes ipv6 threshold 70
!

# Sub interface config
interface xe1.3
 description L3VPN-6VPE-VRF3
 encapsulation dot1q 3
 load-interval 30
 ip vrf forwarding VRF3
 ip address 10.1.23.1/24
 ipv6 address 2001:0:4::1/64
 isis network point-to-point
 isis circuit-type level-1-2
 ip router isis 100
 ipv6 router isis 100
 isis wide-metric 100
 bfd interval 3 minrx 3 multiplier 3
!

# ISIS config
router isis CEPE3 VRF3
 is-type level-1-2
 metric-style wide
 net 49.0001.0000.0000.0003.00
 redistribute bgp
 !
 address-family ipv6
 redistribute bgp
 exit-address-family

# BGP config for VRF
router bgp 1001
 bgp router-id 14.14.14.1
 bgp auto-policy-soft-reset enable
 bgp graceful-restart
 neighbor 12.12.12.1 remote-as 1001
 neighbor 12.12.12.1 update-source 14.14.14.1
 !
 address-family ipv4 unicast
 neighbor 12.12.12.1 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 12.12.12.1 activate
 exit-address-family
 !
 address-family ipv4 vrf VRF3
 redistribute connected
 redistribute isis
 exit-address-family
 !
 address-family ipv6 vrf VRF3
 redistribute connected
 redistribute isis
 exit-address-family
 !
```

© 2026 IP Infusion Inc. Proprietary

Validation for PE2

Validation ensures that the Max VRF FIB Routes Implementation functions as expected by verifying route limits, isolation, and NetConf notification and syslog warning message, preventing resource exhaustion and ensuring efficient routing operations.

```
PE2#show ip bgp vpnv4 all summary
BGP router identifier 14.14.14.1, local AS number 1001
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   Desc
12.12.12.1         4  1001        163        106       2      0      0  00:42:55           116

Total number of neighbors 1

Total number of Established sessions 1

#  Verify VRF Route Threshold (e.g., Max Route Limit)
PE2#show ip route vrf VRF3 summary

------------------------------------------------
IP routing table name is VRF3(4)
------------------------------------------------
IP routing table maximum-paths   : 8
Total number of IPv4 routes      : 10
Total number of IPv4 paths       : 10
Pending routes (due to route max reached): 2
Route Source    Networks
connected       2
isis            8
Total           10
FIB             10

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

VRF FIB Route Limits:
 Configured Route Limit    : 10
 Utilization Percentage    : 100 %
 Action upon reaching limit: stop-install
 Warning Threshold         : 70 %
 Exceeds Threshold         : Yes

PE2#show ip route vrf VRF3 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 "VRF3"
i L2 *>  1.1.1.0/24 [115/11] via 10.1.23.100, xe1.3, installed 00:10:14, last update 00:17:51 ago
i L2 *>  1.1.2.0/24 [115/11] via 10.1.23.100, xe1.3, installed 00:10:14, last update 00:17:51 ago
i L2 *>  1.1.3.0/24 [115/11] via 10.1.23.100, xe1.3, installed 00:10:14, last update 00:17:51 ago
i L2 *>  1.1.4.0/24 [115/11] via 10.1.23.100, xe1.3, installed 00:10:14, last update 00:17:51 ago
i L2 *>  1.1.5.0/24 [115/11] via 10.1.23.100, xe1.3, installed 00:10:14, last update 00:17:51 ago
i L2 *>  1.1.6.0/24 [115/11] via 10.1.23.100, xe1.3, installed 00:10:14, last update 00:17:51 ago
i L2 *>  1.1.7.0/24 [115/11] via 10.1.23.100, xe1.3, installed 00:10:14, last update 00:17:51 ago
i L2 *>  1.1.8.0/24 [115/11] via 10.1.23.100, xe1.3, installed 00:10:14, last update 00:17:51 ago
i L2  >  1.1.9.0/24 [115/11] via 10.1.23.100, xe1.3, installed 00:10:14, last update 00:17:51 ago
```

© 2026 IP Infusion Inc. Proprietary

```
i L2  >  1.1.10.0/24 [115/11] via 10.1.23.100, xe1.3, installed 00:10:14, last update 00:17:51 ago
B     >  10.1.13.0/24 [200/0] via 12.12.12.1, installed 00:10:14, last update 00:39:28 ago
C    *>  10.1.23.0/24 is directly connected, xe1.3, installed 00:10:14, last update 00:33:41 ago
i L1     10.1.23.0/24 [115/10] is directly connected, xe1.3, installed 00:33:41, last update 00:33:41
ago
B     >  10.3.1.1/32 [200/0] via 12.12.12.1, installed 00:10:14, last update 00:39:28 ago
B     >  10.10.2.2/32 [200/0] via 12.12.12.1, installed 00:10:14, last update 00:39:28 ago
C    *>  127.0.0.0/8 is directly connected, lo.VRF3, installed 00:10:14, last update 00:39:56 ago
B     >  192.0.1.0/24 [200/11] via 12.12.12.1, installed 00:10:14, last update 00:39:28 ago
B     >  192.0.2.0/24 [200/11] via 12.12.12.1, installed 00:10:14, last update 00:39:28 ago
B     >  192.0.3.0/24 [200/11] via 12.12.12.1, installed 00:10:14, last update 00:39:28 ago
B     >  192.0.4.0/24 [200/11] via 12.12.12.1, installed 00:10:14, last update 00:39:28 ago
B     >  192.0.5.0/24 [200/11] via 12.12.12.1, installed 00:10:14, last update 00:39:28 ago
B     >  192.0.6.0/24 [200/11] via 12.12.12.1, installed 00:10:14, last update 00:39:28 ago
B     >  192.0.7.0/24 [200/11] via 12.12.12.1, installed 00:10:14, last update 00:39:28 ago
B     >  192.0.8.0/24 [200/11] via 12.12.12.1, installed 00:10:14, last update 00:39:28 ago
B     >  192.0.9.0/24 [200/11] via 12.12.12.1, installed 00:10:14, last update 00:39:28 ago
B     >  192.0.10.0/24 [200/11] via 12.12.12.1, installed 00:10:14, last update 00:39:28 ago

Total number of IPv4 routes 26

Gateway of last resort is not set
PE2#show rib pending-txlist ipv4 vrf VRF3
IPv4 pending:
0x0285ed60: 1.1.9.0/24         req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x0285ef70: 1.1.10.0/24        req(1) new_req(0) fib_notify(0) seq(0) lock(3)

# Ping Validation from PE -- PE
PE2#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         12.12.12.1/32 [115/30] via 31.31.1.2, ce43/4, installed 00:46:19, last update 00:46:19
ago
i L1         13.13.13.1/32 [115/20] via 31.31.1.2, ce43/4, installed 00:37:00, last update 00:37:00
ago
C            14.14.14.1/32 is directly connected, lo, installed 00:46:28, last update 00:46:28 ago
C            31.31.1.0/24 is directly connected, ce43/4, installed 00:46:22, last update 00:46:22 ago
i L1         41.41.1.0/24 [115/20] via 31.31.1.2, ce43/4, installed 00:46:19, last update 00:46:19
ago
C            127.0.0.0/8 is directly connected, lo, installed 00:46:56, last update 00:46:56 ago

Gateway of last resort is not set
PE2#ping 12.12.12.1
Press CTRL+C to exit
PING 12.12.12.1 (12.12.12.1) 100(128) bytes of data.
108 bytes from 12.12.12.1: icmp_seq=1 ttl=64 time=0.490 ms
108 bytes from 12.12.12.1: icmp_seq=2 ttl=64 time=0.478 ms
108 bytes from 12.12.12.1: icmp_seq=3 ttl=64 time=0.454 ms
108 bytes from 12.12.12.1: icmp_seq=4 ttl=64 time=0.791 ms

--- 12.12.12.1 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3065ms
rtt min/avg/max/mdev = 0.454/0.553/0.791/0.137 ms


# Show mpls vrf forwarding output
PE2#show mpls vrf-forwarding-table vrf VRF3
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
(m) - Service mapped over multipath transport
(e) - Service mapped over ECMP
```

```
(D) - Down
Ext-Color - Extended-community color advertised by BGP
B(x) - BGP EVPN MPLS Services

Code    FEC                FTN-ID VRF-ID    Nhlfe-ID    Pri   Out-Label    Out-
Intf         Nexthop         UpTime    Ext-Color
   B>   10.1.13.0/24        1         4      8           -     -            -                -
               00:47:08  -
                                              3           Yes   25602        -
                12.12.12.1      -         -
   B>   10.3.1.1/32         2         4      8           -     -            -                -
               00:47:08  -
                                              3           Yes   25602        -
                12.12.12.1      -         -
   B>   10.10.2.2/32        3         4      8           -     -            -                -
               00:47:08  -
                                              3           Yes   25602        -
                12.12.12.1      -         -
   B>   192.0.1.0/24        4         4      8           -     -            -                -
               00:47:08  -
                                              3           Yes   25602        -
                12.12.12.1      -         -
   B>   192.0.2.0/24        5         4      8           -     -            -                -
               00:47:08  -
                                              3           Yes   25602        -
                12.12.12.1      -         -
   B>   192.0.3.0/24        6         4      8           -     -            -                -
               00:47:08  -
                                              3           Yes   25602        -
                12.12.12.1      -         -
   B>   192.0.4.0/24        7         4      8           -     -            -                -
               00:47:08  -
                                              3           Yes   25602        -
                12.12.12.1      -         -
   B>   192.0.5.0/24        8         4      8           -     -            -                -
               00:47:08  -
                                              3           Yes   25602        -
                12.12.12.1      -         -
   B>   192.0.6.0/24        9         4      8           -     -            -                -
               00:47:08  -
                                              3           Yes   25602        -
                12.12.12.1      -         -
   B>   192.0.7.0/24        10        4      8           -     -            -                -
               00:47:08  -
                                              3           Yes   25602        -
                12.12.12.1      -         -
   B>   192.0.8.0/24        11        4      8           -     -            -                -
               00:47:08  -
                                              3           Yes   25602        -
                12.12.12.1      -         -
   B>   192.0.9.0/24        12        4      8           -     -            -                -
               00:47:08  -
                                              3           Yes   25602        -
                12.12.12.1      -         -
   B>   192.0.10.0/24       13        4      8           -     -            -                -
               00:47:08  -
                                              3           Yes   25602        -
                12.12.12.1      -         -
   B>   2000::/64           2504      4      8           -     -            -                -
               00:47:09  -
                                              3           Yes   25602        -
                12.12.12.1      -         -
   B>   2000::1:0:0:0:0/64  2505      4      8           -     -            -                -
               00:47:09  -
                                              3           Yes   25602        -
                12.12.12.1      -         -
   B>   2000::2:0:0:0:0/64  2506      4      8           -     -            -                -
               00:47:09  -
                                              3           Yes   25602        -
                12.12.12.1      -         -
   B>   2000::3:0:0:0:0/64  2507      4      8           -     -            -                -
```

```
              00:47:09  -
                                          3           Yes   25602       -
               12.12.12.1      -         -
   B>   2000::4:0:0:0:0/64  2508      4      8           -     -           -                -
              00:47:09  -
                                          3           Yes   25602       -
               12.12.12.1      -         -
   B>   2000::5:0:0:0:0/64  2509      4      8           -     -           -                -
              00:47:09  -
                                          3           Yes   25602       -
               12.12.12.1      -         -
   B>   2000::6:0:0:0:0/64  2510      4      8           -     -           -                -
              00:47:09  -
                                          3           Yes   25602       -
               12.12.12.1      -         -
   B>   2000::7:0:0:0:0/64  2511      4      8           -     -           -                -
              00:47:09  -
                                          3           Yes   25602       -
               12.12.12.1      -         -
   B>   2000::8:0:0:0:0/64  2512      4      8           -     -           -                -
              00:47:09  -
                                          3           Yes   25602       -
               12.12.12.1      -         -
   B>   2000::9:0:0:0:0/64  2513      4      8           -     -           -                -
              00:47:09  -
                                          3           Yes   25602       -
               12.12.12.1      -         -
   B>   2001:0:3::/64       4004      4      8           -     -           -                -
              00:47:09  -
                                          3           Yes   25602       -
               12.12.12.1      -         -
   B>   2001:db8:2::2/128   4005      4      8           -     -           -                -
              00:47:09  -
                                          3           Yes   25602       -
               12.12.12.1      -         -
7001-PE2#

# IPv6
PE2#show ipv6 route vrf VRF3 summary

-----------------------------------------------
IPv6 routing table name is VRF3(4)
-----------------------------------------------
IPv6 routing table maximum-paths : 8
Total number of IPv6 routes      : 10
Total number of IPv6 paths       : 10
Pending routes (due to route max reached): 3
Route Source    Networks
connected       3
isis            7
Total           10
FIB             10

ECMP statistics (active in ASIC):
 Total number of IPv6 ECMP routes   : 0
 Total number of IPv6 ECMP paths    : 0

VRF FIB Route Limits:
 Configured Route Limit    : 10
 Utilization Percentage    : 100 %
 Action upon reaching limit: stop-install
 Warning Threshold         : 70 %
 Exceeds Threshold         : Yes

PE2#show rib pending-txlist ipv6 vrf VRF3
IPv6 pending:
0x0285ccc0: 2222::7:0:0:0:0/64 req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x0285ced0: 2222::8:0:0:0:0/64 req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x0285d0e0: 2222::9:0:0:0:0/64 req(1) new_req(0) fib_notify(0) seq(0) lock(3)
```

© 2026 IP Infusion Inc. Proprietary

```
# Ping VPNV4 prefixes
PE2#ping mpls l3vpn VRF3 10.1.13.0/24
Sending 5 MPLS Echos to 10.1.13.0, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

 Type 'Ctrl+C' to abort

 !

Success Rate is 100.00 percent (5/5)

PE2#ping mpls l3vpn VRF3 192.0.1.0/24
Sending 5 MPLS Echos to 192.0.1.0, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

 Type 'Ctrl+C' to abort

 !
 !
 !
 !
 !

Success Rate is 100.00 percent (5/5)

PE2#ping mpls l3vpn VRF3 192.0.5.0/24
Sending 5 MPLS Echos to 192.0.5.0, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

 Type 'Ctrl+C' to abort

 !
 !
 !
 !
 !

Success Rate is 100.00 percent (5/5)
```

Maximum FIB Routes Commands

The VRF FIB introduces the following configuration commands in OcNOS.

© 2026 IP Infusion Inc. Proprietary

maximum-fib-routes threshold

Use this command to set a route limit in a VRF FIB and optionally configure a threshold percentage using the threshold keyword.

Note: This generates a warning when the current route count exceeds the threshold. The threshold can only be configured if a route limit is set.

Use the **no** form of the command with threshold keyword to remove configured threshold percentage for the route limit.

Command Syntax

```
maximum-fib-routes (ipv4|ipv6) threshold <1-100>
no maximum-fib-routes (ipv4|ipv6) threshold
```

Parameters

ipv4|ipv6

Configures the maximum route limit for the specified address family (IPv4 or IPv6).

`threshold<1 - 100>`

Triggers a NetConf notification and syslog warning message when the route count exceeds the specified percentage of the maximum limit.

Default

80%.

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.6.1.

Example

```
(config)#ip vrf VRF_A
(config-vrf)# maximum-fib-routes ipv4 10000 stop-install
(config-vrf)# maximum-fib-routes ipv4 threshold 90
(config)#ip vrf VRF_B
(config-vrf)# maximum-fib-routes ipv4 5000 warning-only
```

© 2026 IP Infusion Inc. Proprietary

maximum-fib-routes

Use this command to set a limit on the number of dynamic routes installed in a VRF FIB, excluding static and connected routes.

Notes:

- The limit applies only to dynamic routes and does not apply to static or connected routes.
- If recursive routes exist and the configured limit is lower than the current route count in the FIB, the resolved route information for recursive routes may be incomplete in the "`show ip route`" command output.
- Execute `clear ip route <vrf-name>` or `clear ipv6 route <vrf-name>` when the FIB route count already exceeds the configured limit.

Use the `no` form of this command without the threshold keyword to disable the maximum routes limit for a VRF.

Command Syntax

```
maximum-fib-routes (ipv4|ipv6) <1 - 2147483647> (stop-install | warning-only)
no maximum-fib-routes (ipv4|ipv6)
```

Parameters

ipv4 <1 - 2147483647>
Specify the IPv4 address family for which to set the maximum route limit in the VRF FIB.

ipv6 <1 - 2147483647>
Specify the IPv6 address family for which to set the maximum route limit in the VRF FIB.

stop-install
Block the installation of additional dynamic routes in the VRF FIB once the configured maximum limit is reached.

warning-only
Defines the maximum routes limit only as an advisory guideline. An advisory limit triggers a warning when the number of routes in the VRF FIB reaches the maximum limit but does not reject additional routes.

Default

Disabled

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.6.1.

Example

```
(config)#ip vrf VRF_A
(config-vrf)# maximum-fib-routes ipv4 10000 stop-install
(config)#ip vrf VRF_B
(config-vrf)# maximum-fib-routes ipv4 5000 warning-only
```

© 2026 IP Infusion Inc. Proprietary

Glossary

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

Table 112. Glossary

Key Terms/Acronym	Description
MPLS	Multiprotocol Label Switching (MPLS): A high-performance network routing technique that directs data using labels rather than network addresses, improving speed and efficiency.
L3VPN	Layer 3 Virtual Private Network (L3VPN): A VPN that operates at the network layer, using MPLS to connect multiple customer sites while maintaining routing separation.
VRF	Virtual Routing and Forwarding (VRF): A technology that enables multiple isolated routing tables on a single router, allowing separate VPN instances.
PE	Provider Edge (PE) - A router at the network edge that connects customer sites to the service provider's MPLS backbone.
CE	Customer Edge (CE): A router at the customer site that connects to the PE router and exchanges routing information.
FIB	Forwarding Information Base (FIB) - A routing table used by routers to forward packets efficiently.
EBGP	External Border Gateway Protocol (EBGP): A routing protocol used for communication between different autonomous systems, commonly used between PE and CE routers.
IGP	Interior Gateway Protocol (IGP) - A routing protocol used within a single network, such as OSPF or IS-IS, to manage routes in the MPLS backbone.

© 2026 IP Infusion Inc. Proprietary

VIRTUAL ROUTER REDUNDANCY PROTOCOL CONFIGURATION

VRRP Configuration 2926
Terminology 2926
VRRP Process 2926
One Virtual Router 2928
Two Virtual Routers 2931
Two Backup Routers 2935
Interface Tracking 2940
VRRP-Backward Compatibility 2947
Redundancy Using VRRP and OSPF: Two Virtual Routers 2950
VRRP Over MLAG 2954
Custom VRF Configuration 2961
Object Tracking Using IP SLA 2966
Overview 2966
Object Tracking Using IP SLA Configuration 2966
VRRP IPv6 Configuration 2973
Topology 2973
Configuration 2973
Disabling the Master 2976
GLOBAL IPV6 SUPPORT FOR VRRP 2977

© 2026 IP Infusion Inc. Proprietary

VRRP Configuration

This section 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.

For complete information about MC-LAG configuration with Intra-domain-Link (IDL) and Intra-domain-peer (IDP), refer to the OcNOS Layer 2 document.

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 202, the problem with this configuration is that it produces a single point of failure if this first-hop router fails.

© 2026 IP Infusion Inc. Proprietary

Figure 202. VRRP Process - First-Hop Not Reachable

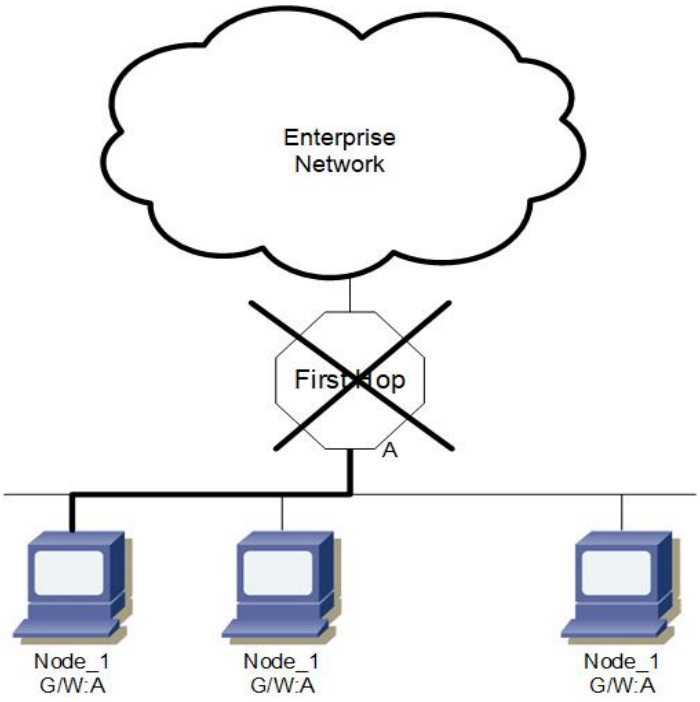

The Virtual Router Redundancy Protocol attempts to solve this problem by introducing the concept of a virtual router, composed of two or more VRRP routers on the same subnet as shown in Figure 203. The concept of a virtual IP address is also introduced, which is the address that end hosts configure as their default gateway. One of the routers called the “Master” forwards packets on behalf of this IP address.

Figure 203. VRRP Process - Master and Backup VR

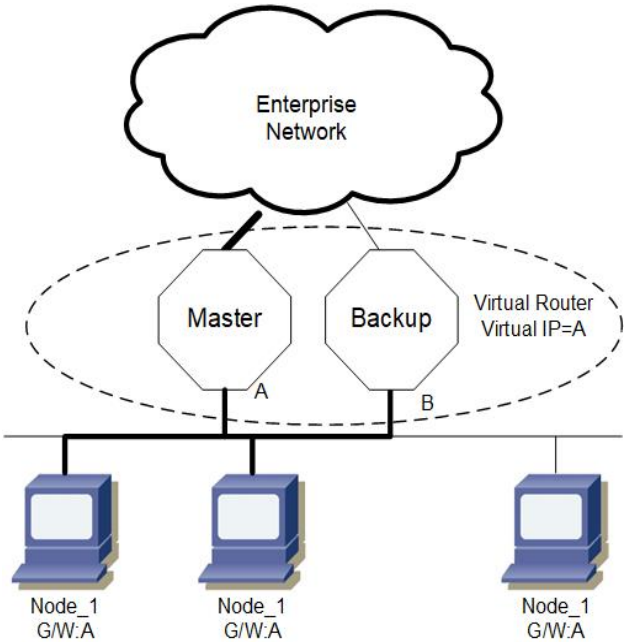

As shown in Figure 204, if the Master router fails, one of the other routers (Backup) assumes forwarding responsibility for it.

© 2026 IP Infusion Inc. Proprietary

Figure 204. VRRP Process - Master Down and Backup Takeover

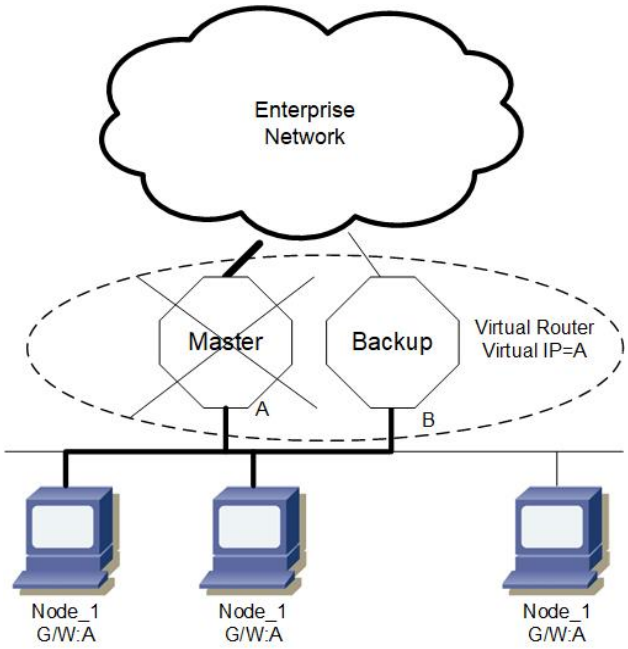

At first glance, the configuration in might not seem very useful, as it doubles the cost, and leaves one router idle at all times. This, however, can be avoided by creating two virtual routers and splitting the traffic between them.

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.

© 2026 IP Infusion Inc. Proprietary

Topology

Figure 205. VRRP with One Virtual Router

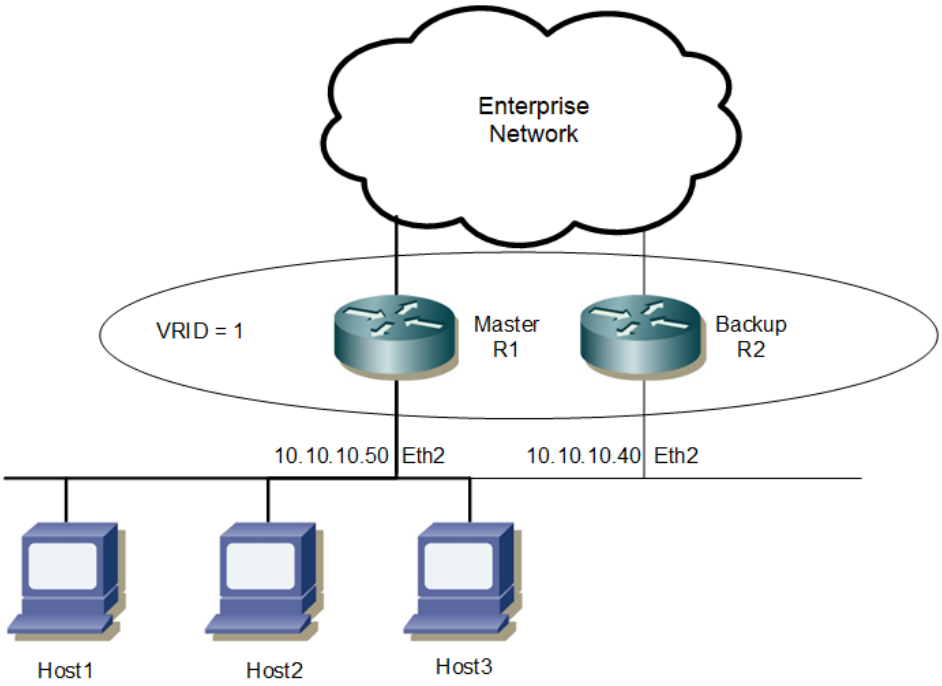

Configuration

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.

© 2026 IP Infusion Inc. Proprietary

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

Router

```
#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 206. Configuring VRRP with Two Virtual Routers

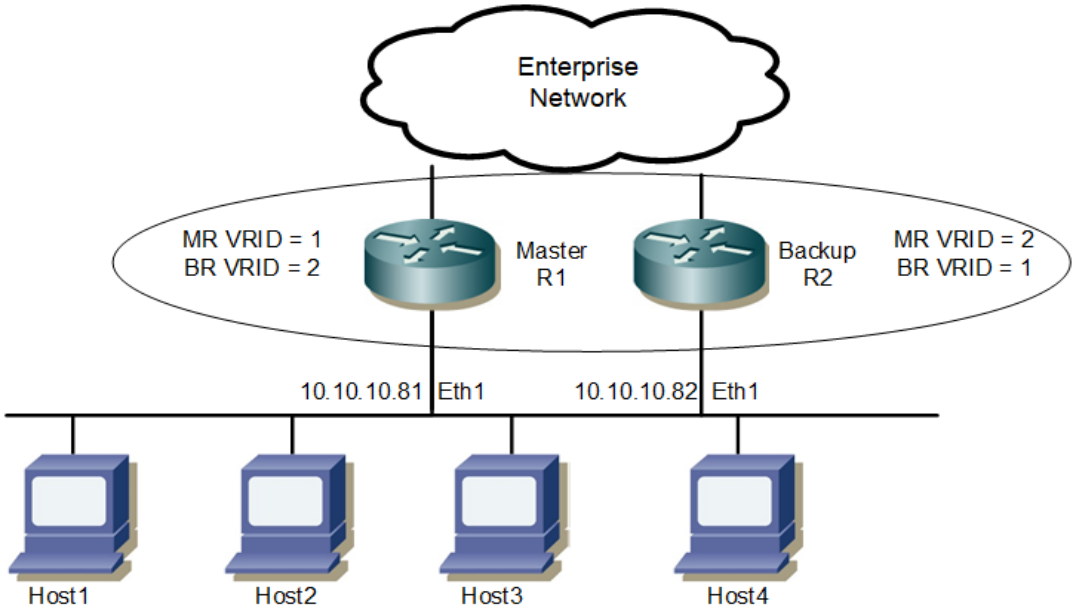

Configuration

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),

© 2026 IP Infusion Inc. Proprietary

	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 rou ter 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 207. Configuring VRRP with Two Backup Routers

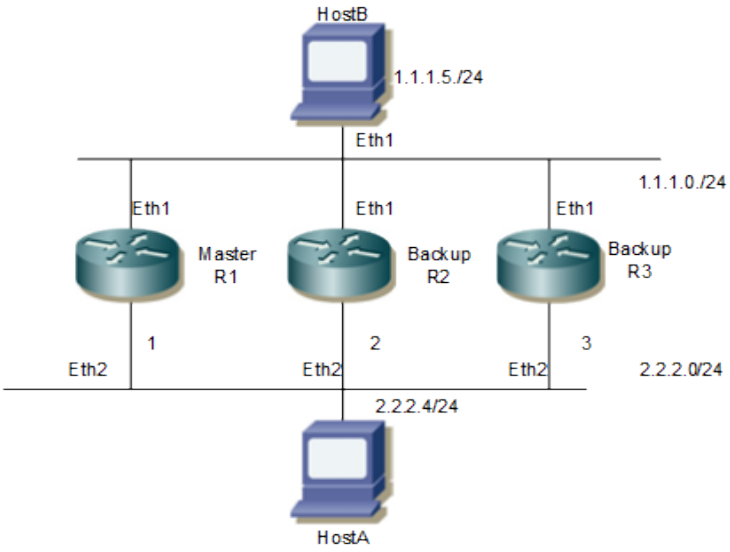

Configuration

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.

© 2026 IP Infusion Inc. Proprietary

`(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.

© 2026 IP Infusion Inc. Proprietary

Validation

R1

```
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
```

R2

```
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
```

R3

```
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
```

© 2026 IP Infusion Inc. Proprietary

```
 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

R1

```
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
```

R3

```
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.

© 2026 IP Infusion Inc. Proprietary

Topology

Figure 208. VRRP Interface Tracking

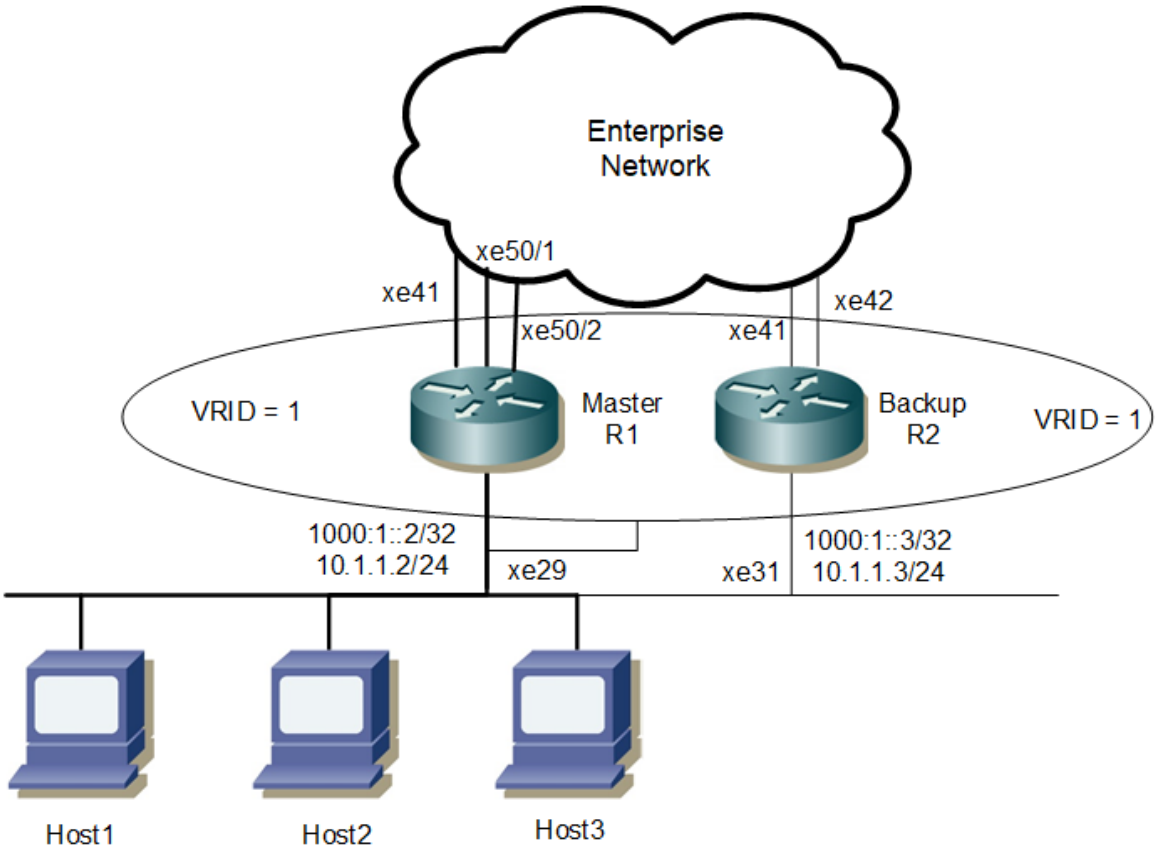

Configuration

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

© 2026 IP Infusion Inc. Proprietary

`(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

© 2026 IP Infusion Inc. Proprietary

(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

© 2026 IP Infusion Inc. Proprietary

(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
```

© 2026 IP Infusion Inc. Proprietary

```
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
```

© 2026 IP Infusion Inc. Proprietary

```
 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

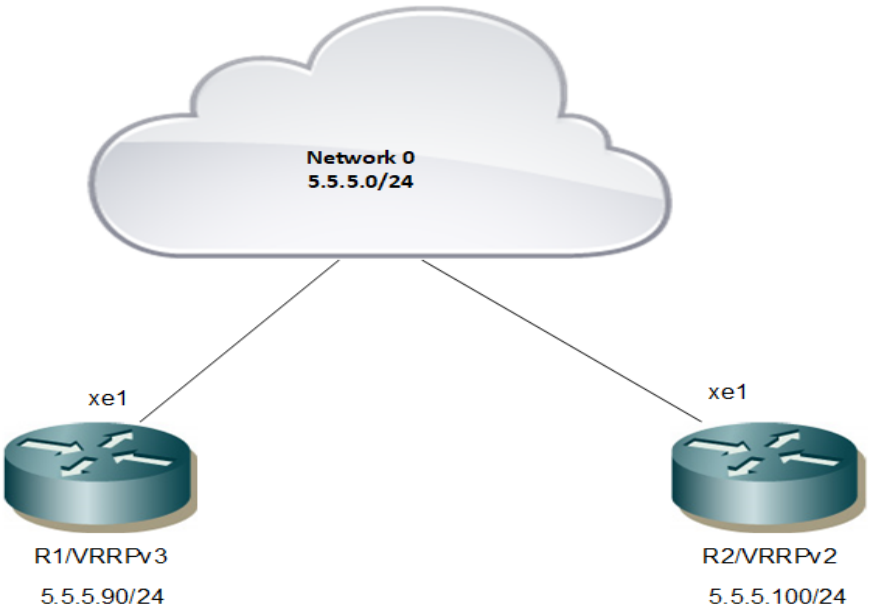

Configuration

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 the router 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 the router.
`(config-router)#commit`	Commit the candidate configuration to the running configuration.

© 2026 IP Infusion Inc. Proprietary

(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.

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

Validation

R1

```
#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

```
#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 209. Redundancy Using VRRP and OSPF

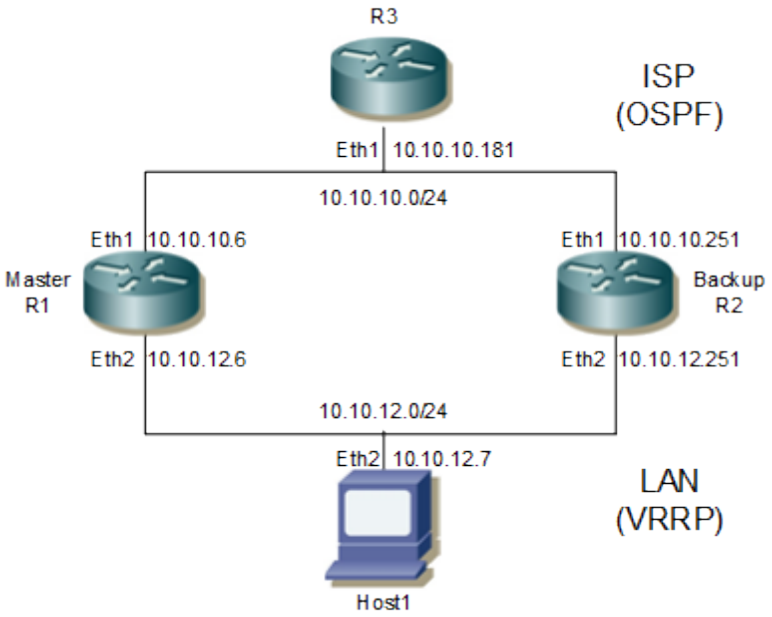

In the above figure:

- 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)
```

© 2026 IP Infusion Inc. Proprietary

```
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.

Configuration

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
```

© 2026 IP Infusion Inc. Proprietary

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 210. VRRP over MLAG

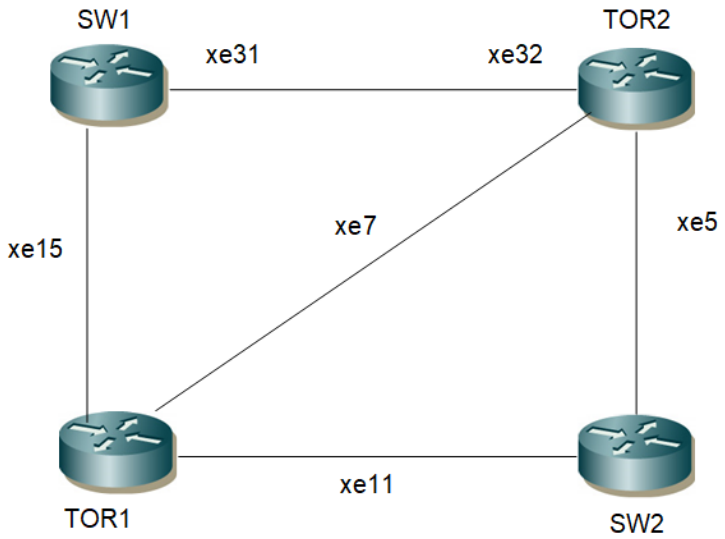

Configuration

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

© 2026 IP Infusion Inc. Proprietary

`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.

© 2026 IP Infusion Inc. Proprietary

`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

© 2026 IP Infusion Inc. Proprietary

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

© 2026 IP Infusion Inc. Proprietary

`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

© 2026 IP Infusion Inc. Proprietary

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

© 2026 IP Infusion Inc. Proprietary

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
```

Custom VRF Configuration

This section shows how to configure VRRP over MLAG with Custom VRF.

Topology

For topology, refer to Figure 210. VRRP over MLAG (page 2954).

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.

© 2026 IP Infusion Inc. Proprietary

`TOR1(config-vrf)#end`	Returns to privilege mode.

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-MPLS-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***
  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
!
```

© 2026 IP Infusion Inc. Proprietary

```
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
!
```

© 2026 IP Infusion Inc. Proprietary

```
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
```

© 2026 IP Infusion Inc. Proprietary

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 Object-Action Mapping Database (OAMD).

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: When both BFD and Track are enabled for a static route, the route is not installed if either BFD or the tracked object is down.

Object Tracking Using IP SLA Configuration

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.

© 2026 IP Infusion Inc. Proprietary

Figure 211. VRRP Object Tracking

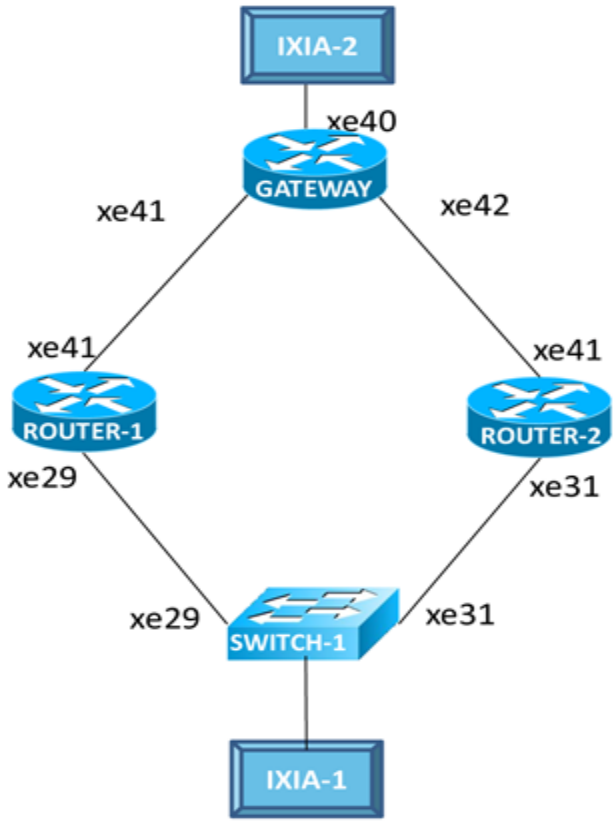

R1 Configuration

`(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.

© 2026 IP Infusion Inc. Proprietary

`(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.

© 2026 IP Infusion Inc. Proprietary

`(config-if)#exit`	Exit interface mode.
`(config)#exit`	Exit configure mode.

R2 Configuration

`(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

© 2026 IP Infusion Inc. Proprietary

`(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  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,  P
riority  Delta:  30,  Status:  UP
Auth-type:  simple  text,  String:  abcd
Multicast  membership  on  IPv4  interface  vlan1.1001:  JOINED  V  2-Compatible:  TRUE

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 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  sec  Accept  mode: TRUE
Preempt  mode:  TRUE
Auth-type:  simple  text,  String:  abcd
Multicast  membership  on  IPv4  interface  vlan1.1001:  JOINED  V2-Compatible:  TRUE

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

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 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  sec  Accept  mode: TRUE
Preempt  mode:  TRUE
Multicast  membership  on  IPv6  interface  vlan1.1001:  JOINED  V2-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  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
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  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

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 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:  JOINED  V2-Compatible:  FALSE

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  Priority 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  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

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 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:  JOINED  V2-Compatible:  FALSE
```

VRRP IPv6 Configuration

This section 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

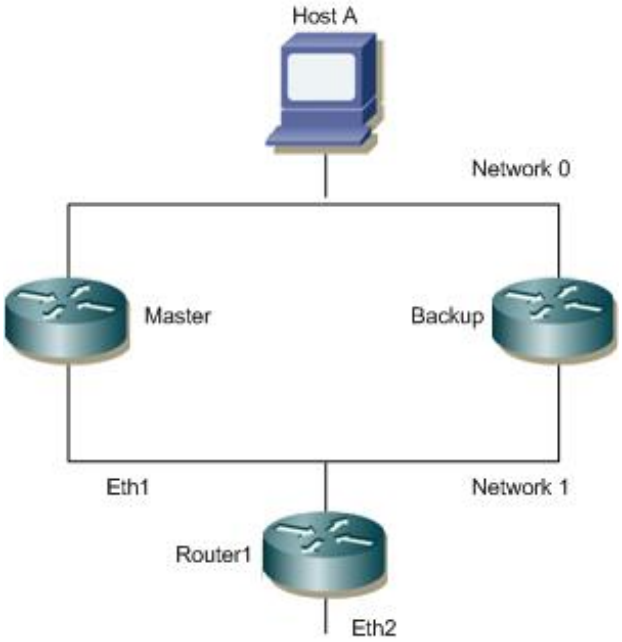

Figure 212. Topology

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.

© 2026 IP Infusion Inc. Proprietary

`(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.

© 2026 IP Infusion Inc. Proprietary

`(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
```

© 2026 IP Infusion Inc. Proprietary

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

© 2026 IP Infusion Inc. Proprietary

(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
```

© 2026 IP Infusion Inc. Proprietary

VIRTUAL ROUTER REDUNDANCY PROTOCOL COMMAND REFERENCE

VRRP Commands 2983
accept-mode 2984
advertisement-interval 2985
authentication text 2986
circuit-failover 2987
debug vrrp 2988
disable 2989
enable 2990
ip pim redundancy 2991
ipv4-exclude-pseudo-header 2992
operational-ip 2993
preempt-mode 2994
priority 2995
router vrrp 2996
show debugging vrrp 2997
show running-config router vrrp 2998
show running-config router ipv6 vrrp 2999
show vrrp 3000
show vrrp (global | ipv4) statistics 3002
show vrrp 3003
show vrrp summary 3004
snmp restart vrrp 3005
switch-back-delay 3006
track decrement 3007
undebug vrrp 3008
virtual-ip 3009
vrrp compatible-v2 3010
vrrp ipv4-exclude-pseudo-header 3011
vrrp vmac 3012
VRRP v6 Commands 3013
advertisement-interval 3014
circuit-failover 3015
disable 3016
enable 3017
preempt-mode 3018

© 2026 IP Infusion Inc. Proprietary

priority 3019
router ipv6 vrrp 3020
router ipv6 vrrp vlan 3021
show running-config vrrpv6 3022
virtual-ipv6 3023

© 2026 IP Infusion Inc. Proprietary

VRRP Commands

This section describes the commands for VRRP.

accept-mode 2984
advertisement-interval 2985
authentication text 2986
circuit-failover 2987
debug vrrp 2988
disable 2989
enable 2990
ip pim redundancy 2991
ipv4-exclude-pseudo-header 2992
operational-ip 2993
preempt-mode 2994
priority 2995
router vrrp 2996
show debugging vrrp 2997
show running-config router vrrp 2998
show running-config router ipv6 vrrp 2999
show vrrp 3000
show vrrp (global | ipv4) statistics 3002
show vrrp 3003
show vrrp summary 3004
snmp restart vrrp 3005
switch-back-delay 3006
track decrement 3007
undebug vrrp 3008
virtual-ip 3009
vrrp compatible-v2 3010
vrrp ipv4-exclude-pseudo-header 3011
vrrp vmac 3012

© 2026 IP Infusion Inc. Proprietary

accept-mode

Use this command to enable/disable accept mode for the session.

Controls whether a VRRP master node will accept/respond to packets addressed to the Virtual-IP address as its own address if it is not the Virtual-IP Owner.

Command Syntax

```
accept-mode true
accept-mode false
```

Parameter

None

Default

Enabled

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

# configure terminal
(config)#router vrrp 2 eth2
(config-router)#accept-mode True
```

© 2026 IP Infusion Inc. Proprietary

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: Higher-priority master routers with slower transmission rates than their backup routers are unstable. This is because low-priority nodes configured to faster rates could come online and decide they should be masters before they have heard anything from the higher-priority master with a slower rate. When this happens, it is temporary: once the lower-priority node does hear from the higher-priority Mmaster, it will relinquish mastership.

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

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
```

© 2026 IP Infusion Inc. Proprietary

authentication text

Use this command to enable authentication of VRRPv2 packets received from other routers in the group.

If configured, all routers within the VRRP group must use the same authentication string.

Authentication is supported only for VRRPv2 packets. V2-compatibility must be enabled.

Use the `no` form of this command to remove text authentication.

Command Syntax

```
authentication text TEXT-STRING
no authentication text
```

Parameter

TEXT-STRING

Password, maximum eight alphanumeric characters

Default

Disabled.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 4.2.

Examples

```
(config)#
(config)#router vrrp 1 eth3
(config-router)#v2-compatible
(config-router)#authentication text abc_123
(config-router)#end
#

(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` parameter with this command to disable this feature.

Command Syntax

```
circuit-failover [IFNAME] |<1-253>
no circuit-failover [IFNAME] |<1-253>
```

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.

Default

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
```

© 2026 IP Infusion Inc. Proprietary

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 execution 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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

ip pim redundancy

Use this command to set the priority for which a router is elected as the designated router (DR).

Note: This command should be applied to the all related VRRP routers with identical priority values.

Command Syntax

```
ip pim redundancy <1-255> vrrp dr-priority <0-4294967294>
no ip pim redundancy vrrp
```

Parameter

<1-255>
VRRP virtual router identifier

<0-4294967294>
DR priority

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS version 4.0.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#ip pim redundancy 1 vrrp dr-priority 900

(config)#interface eth0
(config-if)#no ip pim redundancy vrrp
```

© 2026 IP Infusion Inc. Proprietary

ipv4-exclude-pseudo-header

Use this command to exclude the pseudo-header in IPv4 VRRPv3 checksum calculation on the VRRP group in router mode to support IPv4 VRRPv3 interoperability.

Note: Use this command when checksum errors are observed with other vendors for VRRPv3.

Use the `no` form of this command to delete the VRRP IPv4 checksum pseudo-header exclude configuration from the VRRP group on the router.

Command Syntax

```
ipv4-exclude-pseudo-header
```

Parameter

None

Default

Disabled

Command Mode

Router mode

Applicability

This command was introduced in OcNOS-SP version 2.4.

Examples

```
#configure terminal
(config)#router vrrp 2 eth2
(config-router)#ipv4-exclude-pseudo-header

#configure terminal
(config)#router vrrp 2 eth2
(config-router)#no ipv4-exclude-pseudo-header
```

© 2026 IP Infusion Inc. Proprietary

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

None

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
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

priority

Use this command to configure the priority to use by this VRRP router in master election. The value of 255 (decimal) is reserved for the router that owns the IP 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.

Use the `no` parameter with this command to restore the default setting.

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 50
```

© 2026 IP Infusion Inc. Proprietary

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

None

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router vrrp 1 eth0
```

© 2026 IP Infusion Inc. Proprietary

show debugging vrrp

Use this command to display the set VRRP debugging option.

Command Syntax

```
show debugging vrrp
```

Parameters

None

Command Mode

Execution mode and Privileged execution 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
```

© 2026 IP Infusion Inc. Proprietary

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

Execution 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
!
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution 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
!
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution 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
```

The following table explains the show command 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.

© 2026 IP Infusion Inc. Proprietary

Field	Description
Priority	Configured VRRP priority for the interface.
Advertisement interval	Configured VRRP advertisement interval.
Preempt mode	Whether preemption is allowed on the interface.

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.8.

Example

```
#show vrrp global statistics
VRRP Global Statistics
Checksum Errors  : 0
Version Errors   : 0
VRid Errors      : 0
Discontinuity Time: 00 hour, 00 min, 00 sec


#show 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
```

© 2026 IP Infusion Inc. Proprietary

show vrrp

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

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show vrrp 7 eth0
```

© 2026 IP Infusion Inc. Proprietary

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

Execution mode and Privileged execution mode

Applicability

This command was introduced in OcNOS 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
```

The following table explains the show command output fields.

Table 113. 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.

© 2026 IP Infusion Inc. Proprietary

snmp restart vrrp

Use this command to restart SNMP in Virtual Routing Redundancy Protocol (VRRP).

Command Syntax

```
snmp restart vrrp
```

Parameters

None

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)# snmp restart vrrp
```

© 2026 IP Infusion Inc. Proprietary

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

Zero

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
```

© 2026 IP Infusion Inc. Proprietary

track decrement

Use this command to decrement the priority value during a down event.

Use the no form of this command to not decrement the priority value during a down event.

Command Syntax

```
track <0-500> decrement <1-253>
no track <1-500>
```

Parameters

<1-500>

Identifier of the tracked object.

<1-253>

Priority delta value. The value by which the virtual router decrement its priority value during a down event. Configure this value to be greater than the difference of priorities between the master and backup routers.

Command Mode

VRRP router Mode

Applicability

This command was introduced in OcNOS version 5.1

Example

```
(config)#router vrrp 1 eth1
(config-router)#virtual-ip 11.0.0.5
(config-router)#track 1 decrement 30
```

© 2026 IP Infusion Inc. Proprietary

undebug vrrp

Use this command to disable debugging options for VRRP.

Command Syntax

```
undebug vrrp (all|event|packet [send|recv|])
no undebug 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

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#undebug vrrp all
The example below shows how to disable debug options for VRRP events.
#undebug vrrp events
The example below shows how to disable debug options for VRRP packets sent.
#undebug vrrp packet send
The example below shows how to disable debug options for VRRP packets received.
#undebug vrrp packet recv
```

© 2026 IP Infusion Inc. Proprietary

virtual-ip

Use this command to set the Virtual Internet Protocol (VIP) 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.

Command Syntax

```
virtual-ip A.B.C.D (owner|)
no virtual-ip (owner|)
```

Parameters

A.B.C.D
Virtual IP address of the interface that participates in virtual routing.

owner
IP address as the owner.

Default

None

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
```

© 2026 IP Infusion Inc. Proprietary

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

Enabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#vrrp compatible-v2 enable
```

© 2026 IP Infusion Inc. Proprietary

vrrp ipv4-exclude-pseudo-header

Use this command to excludes the pseudo-header in IPv4 VRRPv3 checksum calculation on the VRRP groups in configuration mode to support IPv4 VRRPv3 interoperability.

Note: Use this command when checksum errors are observed with other vendors for VRRPv3.

Command Syntax

```
vrrp ipv4-exclude-pseudo-header (enable| disable)
```

Parameters

enable
Enable ipv4-exclude-pseudo-header

disable
Disable ipv4-exclude-pseudo-header

Default

By default, VRRP includes the pseudo header in checksum calculation.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS-SP version 2.4.

Examples

```
#configure terminal
(config)#vrrp ipv4-exclude-pseudo-header enable

#configure terminal
(config)#vrrp ipv4-exclude-pseudo-header disable
```

© 2026 IP Infusion Inc. Proprietary

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

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
```

© 2026 IP Infusion Inc. Proprietary

VRRP v6 Commands

This section describes the commands for VRRP IPv6.

advertisement-interval 3014
circuit-failover 3015
disable 3016
enable 3017
preempt-mode 3018
priority 3019
router ipv6 vrrp 3020
router ipv6 vrrp vlan 3021
show running-config vrrpv6 3022
virtual-ipv6 3023

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

disable

Use this command to disable a VRRP session on the router (to stop the router from participating in virtual routing). Refer to enable (page 3017) 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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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)#
```

© 2026 IP Infusion Inc. Proprietary

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)#
```

© 2026 IP Infusion Inc. Proprietary

show running-config vrrpv6

Use this command to show the running configuration for VRRPv6.

Command Syntax

```
show running-config router vrrpv6
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced before OcNOS version 3.0.

Example

```
#show running-config vrrpv6
!
vrrp vmac disable
!
!
router ipv6 vrrp 1 eth1
virtual-ipv6 fe80::c0
enable
!
```

© 2026 IP Infusion Inc. Proprietary

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.

Command Syntax

```
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.

Examples

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
```

© 2026 IP Infusion Inc. Proprietary

UNIDIRECTIONAL LINK DETECTION CONFIGURATION

Unidirectional Link Detection Configuration 3025
Topology 3025
Validation 3026
Enable UDLD under bridge-group 3027

© 2026 IP Infusion Inc. Proprietary

Unidirectional Link Detection Configuration

This section shows a complete configuration to enable UDLD in a simple network topolog y.

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 213 shows the topology of the UDLD configuration.

Figure 213. UDLD Configuration

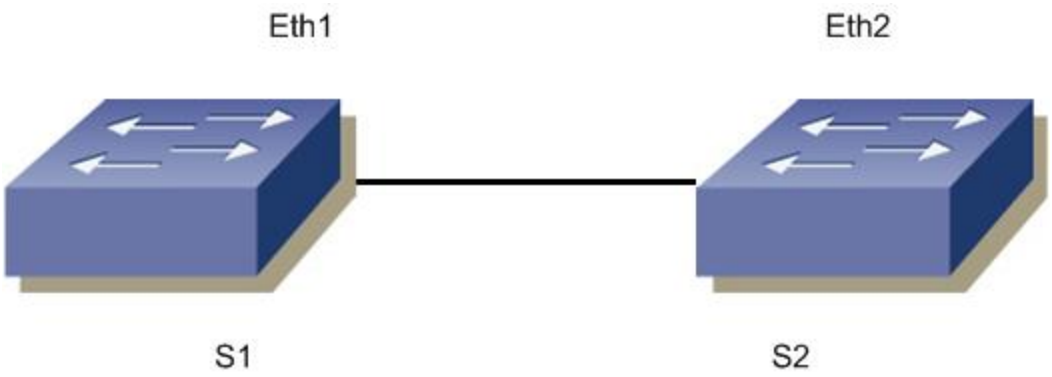

S1

`#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

© 2026 IP Infusion Inc. Proprietary

S2

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

© 2026 IP Infusion Inc. Proprietary

```
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
```

Enable UDLD under bridge-group

S1

`#configure terminal`	Enter configure mode
`(config)#bridge 1 protocol rstp`	Bridge 1 config
`(config)#udld enable`	Enable UDLD globally
`(config)#udld message-time 7`	Configure message time for UDLD packets
`(config-if)#commit`	Commit config.
`(config)#interface eth1`	Enter interface mode
`(config-if)#switchport`	Configure the interface as switch port
`(config-if)#bridge-group 1`	Bridge group 1
`(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)#bridge 1 protocol rstp`	Bridge 1 config.
`(config)#udld enable`	Enable UDLD globally.
`(config)#udld message-time 7`	Configure message time for UDLD packets
`(config-if)#commit`	Commit config.

© 2026 IP Infusion Inc. Proprietary

`(config)#interface eth2`	Enter interface mode
`(config-if)#switchport`	Configure the interface as switch port.
`(config-if)#bridge-group 1`	Bridge group 1
`(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

```
#sh running-config | i bridge 1
bridge 1 protocol rstp

#sh running-config in eth1
 interface eth1
 switchport
 bridge-group 1
 udld state Enable


#sh udld
UDLD                    : Enable
Message Interval(sec)   : 15

Port   UDLD Status    Mode           Link-Status
------------------------------------------------------
eth1    Enable         Normal     Bi-Directional
```

© 2026 IP Infusion Inc. Proprietary

UNIDIRECTIONAL LINK DETECTION COMMAND REFERENCE

Unidirectional Link Detection Commands 3030
udld 3031
udld message-time 3032
udld mode 3033
udld state 3034
show udld 3035
show udld interface 3036

© 2026 IP Infusion Inc. Proprietary

Unidirectional Link Detection Commands

This section describes the Unidirectional Link Detection (UDLD) commands.

udld 3031
udld message-time 3032
udld mode 3033
udld state 3034
show udld 3035
show udld interface 3036

© 2026 IP Infusion Inc. Proprietary

udld

Use this command to enable the UDLD feature globally.

Use `no` form of this command to disable the UDLD feature globally.

Command Syntax

```
udld  enable
no udld enable
```

Parameters

None

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
(config)#udld enable
(config)#no udld enable
```

© 2026 IP Infusion Inc. Proprietary

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
```

© 2026 IP Infusion Inc. Proprietary

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

None

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
```

© 2026 IP Infusion Inc. Proprietary

show udld

Use this command to display UDLD statistic for all interface.

Command Syntax

```
show udld
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#show udld
UDLD                    : Enable
Message Interval(sec)   : 15
Port    UDLD Status     Mode           Link-Status
------------------------------------------------
xe7     Enable          Normal         Bi-Directional
```

Table 114 explains the output fields.

Table 114. 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

© 2026 IP Infusion Inc. Proprietary

show udld interface

Use this command to display UDLD settings for particular interface.

Command Syntax

```
show udld interface IFNAME
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#show udld interface xe14
UDLD Status           : Enable
UDLD Mode             : Aggressive
Link-State            : Bi-Directional
#
```

Table 115 explains the output fields.

Table 115. 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

© 2026 IP Infusion Inc. Proprietary

BRIDGED VIRTUAL INTERFACE

QoS and ACL over Bridge Virtual Interface (BVI) 3038
Overview 3038
Configuration 3039
Running Configuration 3041
Validation 3042
Glossary 3044

© 2026 IP Infusion Inc. Proprietary

QoS and ACL over Bridge Virtual Interface (BVI)

Overview

A Bridge Virtual Interface (BVI) acts as a Layer 3 gateway for a bridged Layer 2 domain, such as a VLAN or a bridge domain. Applying Quality of Service (QoS) and Access Control Lists (ACLs) on a BVI interface allows for intelligent traffic control and filtering at the point where Layer 2 traffic transitions to Layer 3 routing.

Refer to Multi-Protocol Label Switching (MPLS) document for details:

- For BVI Over L3VPN, see the Bridge Virtual Interface (BVI) Over L3VPN section
- For BVI Over 6vPE, see the Bridge Virtual Interface (BVI) Over 6vPE section

Limitations

- Only Q1 and Q2-based platforms are supported.
- L2 subinterfaces support only pop and pop2-tag VLAN rewrite translations and for the purpose of bridge-domain support, all member L2 subinterfaces should have relevant rewrite configurations to make sure that all traffic for bridge-domain has uniform encapsulation.
- The dot1q and dot1ad encapsulations with range are not supported for the BVI.
- BVI is a generic L3 interface and allows users to configure any / all existing OcNOS features. However, only a few of these are supported in release 6.6.0 and those are IP address related, VRF2, any routing protocol specific commands (OSPF/IS-IS/BGP related), MTU, DHCP server/relay, Ingress/Egress ACL and QoS commands related to marking.
- L3 ACL applied at BVI is only relevant for routed traffic. Bridged traffic between L2 sub-interfaces will not be subjected to L3 ACL configured at BVI.
- No Interface counters are supported for BVI Interface.
- BVI can’t be used as a network interface for all transports in MPLS core network.
- For QoS, only marking, policing and rate limiting are supported. Other queuing features are not supported at BVI level. However, existing QoS scheduling and queuing features on other interfaces can be used in conjunction with the BVI interface.
- The following restrictions apply to Q1 based systems:
  - For L3 ACL for BVI only hardware profile ingress-ipv4-subif-ext is supported.
  - For L3 QOS for BVI only hardware profile qos-subif-ext and qos-subif-policer is supported.
  - L2 Mac ACLs are not supported for the BVI.
  - More than one L2 sub interfaces belonging to the same physical interface cannot be part of the same bridge-domain if BVI is configured for the bridge-domain.
- QoS features like shaping, queueing, policing, re-marking are not supported on the BVI interface associated to VPLS.
- BVI interface internally acts as a virtual L3 interface and ACL, QOS will use the same set of cli commands as applicable for a normal L3 interface. However because of some hardware limitations, not all acl and qos features can be supported in the box. Specific hardware profiles need to be enabled to support acl and qos features on bvi interface.

© 2026 IP Infusion Inc. Proprietary

Prerequisites

- After configuring the BVI, ensure that the end-to-end traffic flows without any loss.

Configuration

Topology

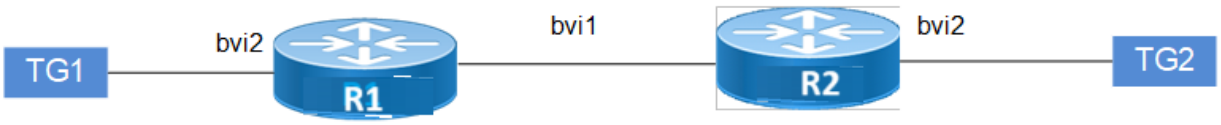

Figure 214. QoS and ACL on BVI Interface in Topology

Configure QoS and ACL on BVI Interface in R1 Router

1. Configure the L2 Subinterface:

```
R1 (config-if)#interface xe19.100 switchport
R1 (config-if)#encapsulation dot1q 100
R1 (config-if)#rewrite pop
R1 (config-if)#commit
R1 (config-if)#exit
```

2. Configure the BVI Interface:

```
R1 (config)#interface bvi1
R1 (config-if)#ip address 172.16.10.1/24
R1 (config-if)#commit
R1 (config-if)#exit
```

3. Map the BVI interface and subinterface into bridge-domain:

```
R1 (config)#bridge-domain 1
R1 (config-bridge-domain )#interface xe19.100
R1 (config-bridge-domain )#routed-interface bvi1
R1 (config-bridge-domain )#commit
R1 (config-bridge-domain )#exit
```

4. Enable QoS and Define Policy:

```
R1(config)#qos enable
R1(config)#qos statistics
R1(config)#commit
R1(config)#exit
R1(config)#class-map type qos match-any C2
R1(config-qos-match-any)#match dscp af11
R1(config-qos-match-any)#match cos 3
R1(config-qos-match-any)#commit
R1(config-qos-match-any)#exit
R1(config)#policy-map type qos P2
R1(config-pmap-qos)#class type qos C2
R1(config-pmap-qos)#police cir 200 mbps eir 200 mbps
R1(config-pmap-qos)#commit
R1(config-pmap-qos)#exit
```

5. ACL configuration:

```
R1(config)#ip access-list test1
R1(config-ip-acl)#10 deny any host 172.16.10.2 host 172.16.10.1
R1(config-ip-acl)#20 permit any any any
R1(config-ip-acl)#commit
R1(config-ip-acl)#exit
```

© 2026 IP Infusion Inc. Proprietary

6. Apply the ACL:

```
R1 (config)#interface bvi1
R1 (config-if)#ip access-group test1 in
R1 (config-if)#commit
R1 (config-if)#exit
```

7. Apply the QoS:

```
R1 (config)#interface bvi1
R1 (config-if)#service-policy type qos input P2
R1 (config-if)#commit
R1 (config-if)#exit
```

Configure QoS and ACL on BVI Interface in R2 Router

1. Configure the L2 Subinterface:

```
R2 (config-if)#interface xe19.100 switchport
R2 (config-if)#encapsulation dot1q 100
R2 (config-if)#rewrite pop
R2 (config-if)#load-interval 30
R2 (config-if)#commit
R2 (config-if)#exit
```

2. Configure the BVI Interface:

```
R2 (config)#interface bvi1
R2 (config-if)#ip address 172.16.10.2/24
R2 (config-if)#commit
R2 (config-if)#exit
```

3. Map the BVI interface and subinterface into bridge-domain:

```
R2 (config)#bridge-domain 1
R2 (config-bridge-domain )#interface xe19.100
R2 (config-bridge-domain )#routed-interface bvi1
R2 (config-bridge-domain )#commit
R2 (config-bridge-domain )#exit
```

4. Enable QoS and Define Policy:

```
R2(config)#qos enable
R2(config)#qos statistics
R2(config)#commit

R2(config)#class-map type qos match-any C2
R2(config-qos-match-any)#match dscp af11
R2(config-qos-match-any)#match cos 3
R2(config-qos-match-any)#commit
R2(config-qos-match-any)#exit

R2(config)#policy-map type qos P2
R2(config-pmap-qos)#police cir 50 mbps eir 50 mbps
R2(config-pmap-qos)#commit
R2(config-pmap-qos)#exit
```

5. ACL configuration:

```
R2(config)#ip access-list test1
R2(config-ip-acl)#10 deny any host 172.16.10.1 host 172.16.10.2
R2(config-ip-acl)#20 permit any any any
R2(config-ip-acl)#commit
R2(config-ip-acl)#exit
```

6. Apply the ACL:

```
R2 (config)#interface bvi1
```

© 2026 IP Infusion Inc. Proprietary

```
R2 (config-if)#ip access-group test1 in
R2 (config-if)#commit
R2 (config-if)#exit
```

7. Apply the QoS:

```
R2 (config)#interface bvi1
R2 (config-if)#service-policy type qos input P2
R2 (config-if)#commit
R2 (config-if)#exit
```

Running Configuration

QoS and ACL on BVI Interface in R1

```
!
interface xe19.100 switchport
 encapsulation dot1q 100
 rewrite pop
 load-interval 30
!
!
interface bvi1
 ip address 172.16.10.1/24
 ip access-group test1 in
!
!
interface bvi1
 ip address 172.16.10.1/24
  service-policy type qos input P2
!
!
bridge-domain 1
 interface xe19.100
 routed-interface bvi1
!
ip access-list test
 10 deny any host 172.16.10.2 host 172.16.10.1
 20 permit any any any
!
qos enable
qos statistics
!
class-map type qos match-any C2
 match dscp af11
 match cos 3

!
policy-map type qos P2
 class type qos C2
  police cir 200 mbps eir 200 mbps
  exit
!
interface bvi1
 service-policy type qos input P2
!
```

QoS and ACL on BVI Interface in R2

```
!
interface xe19.100 switchport
 encapsulation dot1q 100
 rewrite pop
```

```
 load-interval 30
!
bridge-domain 1
 interface xe19.100
 routed-interface bvi1
!
ip access-list test
 10 deny any host 172.16.10.1 host 172.16.10.2
 20 permit any any any
!
qos enable
qos statistics
!
class-map type qos match-any C2
 match dscp af11
 match cos 3
 !
 policy-map type qos P2
 class type qos C2
  police cir 50 mbps eir 50 mbps
  exit
!
interface bvi1
 service-policy type qos input P2
!
interface bvi1
 ip address 172.16.10.2/24
 ip access-group test1 in
!
interface bvi1
 ip address 172.16.10.2/24
  service-policy type qos input P2
!
```

Validation

R1

```
Quality of Service (QoS):


R1#show policy-map interface bvi1

Interface bvi1

Type QoS statistics status : enabled

Type QoS Ingress policy-map : P2

Service-policy (qos) input: P2
---------------------------
Class-map (qos): C2 (match any)
 match dscp af11
 match cos 3
 police cir 200 mbps eir 200 mbps bc 125 kbytes be 125 kbytes
      matched      : 54623616 packets, 6991822848 bytes
      transmitted  : 25285308 packets, 3236519424 bytes
      dropped      : 29338308 packets, 3755303424 bytes


R1#
R1#show interface counters rate mbps
+------------------+-------------+-------------+-------------+-------------+
|     Interface    |    Rx mbps  |    Rx pps   |    Tx mbps  |    Tx pps   |
```

```
+-------------------+--------------+-------------+--------------+-------------+
 ce52                0.00           4             0.01           8
 sa2000              0.00           4             399.36         390000
 xe3                 862.84         842620        0.02           15
 xe19                0.01           8             0.01           12
 xe20                0.00           4             399.39         390027
R1#

 Access Control List (ACL):


R1#show access-lists
ip access-list test1
 1 permit 23 host 172.16.10.2 host 172.16.10.1 [match=58384779]
 20 deny any any any

R1#show interface counters rate mbps
+-------------------+--------------+-------------+--------------+-------------+
|     Interface     |    Rx mbps   |    Rx pps   |    Tx mbps   |    Tx pps   |
+-------------------+--------------+-------------+--------------+-------------+
 ce52                0.00           4             0.01           8
 sa2000              0.00           4             0.01           8
 xe3                 862.86         842633        0.02           15
 xe19                0.01           8             0.01           12
 xe20                0.00           4             0.01           8
R1#
```

R2

```
Quality of Service (QoS):

R2#show policy-map interface bvi1

Interface bvi1

Type QoS Egress policy-map : P2

Service-policy (qos) output: P2
----------------------------
Class-map (qos): C2 (match any)
 match dscp af11
 match cos 3
 police cir 50 mbps eir 50 mbps bc 31 kbytes be 31 kbytes
      matched      : 2374198640 packets, 294400631360 bytes
      transmitted  : 667197196 packets, 82732452304 bytes
    green (output): 41366165640 bytes
   yellow (output): 41366286664 bytes
        red (drop): 211668179056 bytes


R2#
R2#show interface counters rate mbps
+-------------------+--------------+-------------+--------------+-------------+
|     Interface     |    Rx mbps   |    Rx pps   |    Tx mbps   |    Tx pps   |
+-------------------+--------------+-------------+--------------+-------------+
 xe10                0.00           1             100.52         8379
 xe10.20             0.00           0             0.00           1
 xe10.100            0.00           0             100.52         8379
 xe19                4854.29        404532        0.01           8
 xe19.100            4871.83        410921        0.00           3
 xe20                0.01           8             0.01           8

Access Control List (ACL):

R2#show access-lists
```

```
ip access-list test1
 1 permit 23 host 172.16.10.1 host 172.16.10.2 [match=58384779]
 20 deny any any any

 R2#show interface counters rate mbps
+-------------------+--------------+-------------+--------------+-------------+
|     Interface     |    Rx mbps   |    Rx pps   |    Tx mbps   |    Tx pps   |
+-------------------+--------------+-------------+--------------+-------------+
 ce6                 0.01           8             0.00           4
 ce6.6               0.01           4             0.00           4
 sa2000              233.03         227566        0.00           4
 sa2000.20           21.34          24243         0.00           4
 xe10                0.00           1             0.01           3
 xe10.20             0.00           0             0.00           1
 xe10.100            0.00           1             0.01           1
 xe19                0.01           9             0.00           5
 xe19.100            0.01           4             0.00           5
 xe20                233.03         227566        0.00           4
R2#
```

Glossary

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

Key Terms/Acronym	Description
ACL	Access Control List
BVI	Bridge Virtual Interface
DSCP	Differentiated Services Code Point
LDP	Label Distribution Protocol
MTU	Maximum Transmission Unit
QoS	Quality of Service

© 2026 IP Infusion Inc. Proprietary