AOS-CX 10.17 IP Routing Guide

6300, 6400, 8100, 8320, 83xx, 8360, 93xx,

10000 Switch Series

Hewlett Packard Enterprise

Published: November 2025
Version: 1

Copyright Information

© Copyright 2025 Hewlett Packard Enterprise Development LP.

This product includes code licensed under certain open source licenses which require source compliance. The corresponding source for these components is available upon request. This offer is valid to anyone in receipt of this information and shall expire three years following the date of the final distribution of this product version by Hewlett Packard Enterprise Company. To obtain such source code, please check if the code is available in the HPE Software Center at https://myenterpriselicense.hpe.com/cwp-ui/software but, if not, send a written request for specific software version and product for which you want the open source code. Along with the request, please send a check or money order in the amount of US $10.00 to:

Hewlett Packard Enterprise Company
Attn: General Counsel
WW Corporate Headquarters
1701 E Mossy Oaks Rd Spring, TX 77389
United States of America.

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel®, Itanium®, Optane™, Pentium®, Xeon®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the U.S. and other countries.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

All third-party marks are property of their respective owners.

Contents

About this document **18**
- Applicable products 18
- Latest version available online 18
- Command syntax notation conventions 18
- About the examples 19
- Identifying switch ports and interfaces 20
- Identifying modular switch components 20

Virtual Routing and Forwarding (VRF) **22**
- Adding or deleting a VRF 22
- IPv4 static route addition or deletion in a VRF 22
- IPv6 static route addition or deletion in a VRF 23
- Attaching or detaching a port from a VRF 24
- Client applications running on destination VRF 24
- Viewing VRF information 24
- An example of the VRF information provided by the show running-config command 25
- VRF commands 25
  - ip route vrf 25
  - ipv6 route gc interval 27
  - ipv6 route vrf 27
  - show ip route 29
  - show ipv6 route 33
  - show vrf 34
  - vrf 35
  - vrf attach 36

Loopback **38**
- Loopback commands 38
  - interface loopback 38
  - ip address 39
  - ipv6 address 39
  - vrf attach 40
  - show interface loopback 41

Static routing **43**
- Default route 44
- Recursive static routes 44
- Route manager 45
- Configuration concepts 45
- Configuration example procedure 45
  - Basic static route configuration example 45
- Static routing commands 46
  - ip route 46
  - ip route bfd 48
  - ip route vrf 50
  - ipv6 route 51
  - show ip rib 53
  - show ipv6 rib 56
  - ipv6 route vrf 59

show ip route 60
show ip rib host-routes 64
show ipv6 route 65
show ipv6 rib host-routes 66

Open Shortest Path First version 2 (OSPFv2) 68

Overview 68
How OSPFv2 protocol works 69
OSPFv2 concepts 69
OSPFv2 Link-state advertisement (LSA) types 69
OSPFv2 router types 70
OSPFv2 area types 70
OSPFv2 configuration task list 71
Tasks at a glance 72
Configuring OSPF on the routing switch 72
Assigning the routing switch to an OSPF area 73
Setting OSPF network for the area 73
Creating an OSPF virtual link for an area 74
Configuring external route redistribution and control 74
Configuring area ranges on an ABR to reduce advertisements to the backbone 75
Influencing route choice by changing the administrative distance 76
Configuring graceful restart of OSPF routing 76
Configuring OSPF interface settings 76
Configuring OSPF interface authentication 77
Configuring OSPF virtual link settings 78
Prerequisites 78
Configuring OSPF authentication on a virtual link 78
Configuring all OSPF interfaces as passive 79
Configuring SPF throttling timers 80
Viewing OSPFv2 information 80
Clearing OSPF statistics on a switch 80
An example of the OSPFv2 information in the show running-config command 81
OSPFv2 commands 82
active-backbone 82
area (ospf) 83
area default-metric 84
area nssa 85
area range 86
area stub 88
area virtual-link 89
authentication 90
authentication-key 91
clear ip ospf neighbors 93
clear ip ospf statistics 94
dead-interval 95
default-information originate 96
default-information originate always 97
default-metric 98
disable 99
distance 100
distribute-list prefix 102
enable 103
graceful-restart 104
hello-interval 105
ip ospf area 106
ip ospf authentication 108

ip ospf authentication-key 109
ip ospf cost 110
ip ospf dead-interval 111
ip ospf hello-interval 112
ip ospf keychain 113
ip ospf message-digest-key md5 114
ip ospf network 116
ip ospf passive 117
ip ospf priority 118
ip ospf retransmit-interval 118
ip ospf sha-key sha 119
ip ospf shutdown 121
ip ospf transit-delay 122
keychain 123
max-metric router-lsa 124
maximum-paths 125
message-digest-key md5 126
passive-interface default 127
redistribute 128
reference-bandwidth 130
retransmit-interval 131
rfc1583-compatibility 132
router ospf 133
router-id 134
sha-key sha 135
show ip ospf 137
show ip ospf border-routers 138
show ip ospf interface 139
show ip ospf lsdb 142
show ip ospf neighbors 147
show ip ospf routes 149
show ip ospf statistics 153
show ip ospf statistics interface 154
show ip ospf virtual-links 155
summary-address 157
timers lsa-arrival 158
timers throttle lsa 159
timers throttle spf 160
transit-delay 162
trap-enable 163

Open Shortest Path First version 3 (OSPFv3) 164

Overview 164
How OSPFv3 protocol works 165
OSPFv3 protocol 165
OSPFv3 concepts 165
OSPFv3 Link-state advertisement (LSA) types 165
OSPFv3 area types 167
OSPFv3 router types 168
OSPFv3 configuration task list 169
Tasks at a glance 169
Configuring OSPFv3 on the routing switch 170
Creating an OSPFv3 area 170
Setting OSPFv3 network for the area 171
Configuring external route redistribution and control 172
Configuring area ranges on an ABR to reduce advertisements to the backbone 173

Prerequisites ... 173
Procedure ... 173
Influencing route choice by changing the administrative distance ... 173
Configuring graceful restart ... 174
Configuring OSPFv3 virtual link settings ... 174
Configuring OSPFv3 interface settings ... 175
Configuring BFD for OSPFv3 ... 176
Examples ... 176
Configuring all OSPFv3 interfaces as passive ... 177
Configuring SPF throttling timers ... 177
Viewing OSPFv3 information ... 177
Clearing OSPFv3 statistics on a switch ... 178
OSPFv3 commands ... 178
active-backbone ... 178
area ... 179
area authentication ipsec ... 180
area encryption ipsec ... 181
area nssa ... 183
area range ... 185
area stub ... 186
area virtual-link ... 187
area default-metric ... 188
authentication ipsec ... 189
clear ipv6 ospfv3 neighbors ... 191
clear ipv6 ospfv3 statistics ... 192
clear ospfv3-af ipv4 neighbor ... 193
clear ospfv3-af ipv4 statistics ... 194
clear ospfv3-af ipv6 neighbor ... 195
clear ospfv3-af ipv6 statistics ... 197
dead-interval ... 198
default-metric ... 199
disable ... 200
distance ... 201
distribute-list prefix ... 202
enable ... 204
encryption ipsec ... 204
default-information originate ... 207
default-information originate always ... 208
graceful-restart ... 209
hello-interval ... 211
ipv6 ospfv3 area ... 212
ipv6 ospfv3 authentication null ... 213
ipv6 ospfv3 authentication ipsec ... 214
ipv6 ospfv3 cost ... 215
ipv6 ospfv3 dead-interval ... 216
ipv6 ospfv3 encryption ipsec ... 217
ipv6 ospfv3 encryption null ... 219
ipv6 ospfv3 hello-interval ... 220
ipv6 ospfv3 network ... 221
ipv6 ospfv3 passive ... 222
ipv6 ospfv3 priority ... 223
ipv6 ospfv3-af retransmit-interval ... 224
ipv6 ospfv3 shutdown ... 225
ipv6 ospfv3 transit-delay ... 226
maximum-paths ... 227
max-metric router-lsa ... 228

passive-interface default 229
redistribute 230
reference-bandwidth 232
retransmit-interval 233
router-id 234
router ospfv3 235
show ipv6 ospfv3 236
show ipv6 ospfv3 border-routers 238
show ipv6 ospfv3 interface 239
show ipv6 ospfv3 lsdb 241
show ipv6 ospfv3 neighbors 246
show ipv6 ospfv3 routes 248
show ipv6 ospfv3 statistics 250
show ipv6 ospfv3 statistics interface 251
show ipv6 ospfv3 virtual-links 253
summary-address 254
timers lsa-arrival 255
timers throttle lsa 256
timers throttle spf 258
transit-delay 259
trap-enable 260

Open Shortest Path First version 3 Address Family (OSPFv3 AF) 261

Overview 261
Limitations 262
How OSPFv3 AF protocol works 263
OSPFv3 AF protocol 263
OSPFv3 AF concepts 263
OSPFv3 AF Link-state advertisement (LSA) types 263
OSPFv3 AF area types 265
OSPFv3 AF router types 266
OSPFv3 AF configuration task list 267
Tasks at a glance 267
Configuring OSPFv3 AF on the routing switch 267
Assigning the routing switch to an OSPFv3 AF area 268
Setting OSPFv3 AF network for the area 269
Configuring external route redistribution and control 270
Configuring area ranges on an ABR to reduce advertisements to the backbone 271
Prerequisites 271
Procedure 271
Influencing route choice by changing the administrative distance 271
Configuring graceful restart 272
Configuring OSPFv3 AF virtual link settings 272
Configuring OSPFv3 AF interface settings 273
Configuring BFD for OSPFv3 AF 274
Examples 274
Configuring all OSPFv3 AF interfaces as passive 275
Configuring SPF throttling timers 275
Viewing OSPFv3 AF information 276
Clearing OSPFv3 AF statistics on a switch 276
OSPFv3 AF commands 276
active-backbone stub-default-rout 276
address-family unicast 277
area default-metric 278
area nssa 280
area 281

area range 282
area stub 284
area virtual-link 286
authentication ipsec spi 287
clear ospfv3-af ipv4 neighbor 288
clear ospfv3-af ipv4 statistics 290
clear ospfv3-af ipv6 neighbor 291
clear ospfv3-af ipv6 statistics 292
dead-interval 293
debug ospfv3-af 294
default-information originate always 296
default-information originate 297
default-metric 298
diag ospfv3 lsdb dump 299
disable 301
distance 302
distribute-list prefix 303
enable 305
encryption ipsec 305
graceful-restart ignore-lost-interface 308
graceful-restart 309
hello-interval 310
ospfv3-af ipv4 retransmit-interval 311
ospfv3-af ipv4 cost 312
ospfv3-af ipv4 hello-interval 313
ospfv3-af ipv4 network 314
ospfv3-af ipv4 priority 315
ospfv3-af ipv4 shutdown 316
ospfv3-af ipv4 transit-delay 317
ospfv3-af ipv6 area 318
ospfv3-af ipv6 cost 319
ospfv3-af ipv6 dead-interval 320
ospfv3-af ipv6 hello-interval 321
ospfv3-af ipv6 network 322
ospfv3-af ipv6 passive 323
ospfv3-af ipv6 priority 324
ospfv3-af ipv6 retransmit-interval 325
ospfv3-af ipv6 shutdown 326
ospfv3-af ipv6 transit-delay 327
max-metric router-lsa 328
maximum-paths 329
ospfv3-af encryption ipsec spi 330
ospfv3-af authentication ipsec spi 332
ospfv3-af encryption ipsec 334
ospfv3-af ipv4 area 336
ospfv3-af ipv4 dead-interval 337
ospfv3-af authentication null 338
ospfv3-af encryption null 339
ospfv3-af ipv4 passive 339
passive-interface default 340
redistribute 341
reference-bandwidth 344
retransmit-interval 345
router-id 346
router ospfv3-af 347
show ospfv3-af ipv4 statistics 348

show debug buffer module ospfv3-af 349
show ospfv3-af ipv4 353
show ospfv3-af ipv4 border-routers 354
show ospfv3-af ipv4 interface 355
show ospfv3-af ipv4 lsdb 359
show ospfv3-af ipv4 neighbors 365
show ospfv3-af ipv4 routes 368
show ospfv3-af ipv4 statistics interface 370
show ospfv3-af ipv6 373
show ospfv3-af ipv6 border-routers 375
show ospfv3-af ipv6 interface 376
show ospfv3-af ipv6 lsdb 379
show ospfv3-af ipv6 neighbors 385
show ospfv3-af ipv6 routes 388
show ospfv3-af ipv6 statistics 390
show ospfv3-af ipv6 statistics interface 392
show ospfv3-af ipv6 virtual-links 394
show running-config 396
show running-config all 401
show tech ospfv3-af 410
summary-address 417
timers lsa-arrival 418
timers throttle lsa 419
timers throttle spf 420
transit-delay 422

Border Gateway Protocol (BGP) 424

Overview 424
Autonomous system numbers 424
BGP sessions 424
Inter-router communication 424
BGP messages 425
BGP neighbor states 425
Injecting routes/prefixes into the BGP table 426
Path attributes 426
BGP best-path calculation 427
Loop prevention 428
Route policies 428
Resetting BGP sessions 428
IBGP full mesh requirement 429
Route reflectors 429
Loop prevention in route reflectors 430
BGP unnumbered 430
BGP peer groups 430
BGP communities 430
Aggregate routes 430
BGP Graceful-Restart and high availability 431
Configuration tasks list 431
Configuring BGP peers 432
Defining BGP Neighbors 432
Using Loopback Interfaces for BGP peering 432
Verifying BGP routing table 432
Verifying BGP peers 433
Clearing BGP neighbors 434
Address families 434
Scale limits 434

Configuration tasks list 435
Configuring BGP peers 435
Defining BGP Neighbors 435
Using Loopback Interfaces for BGP peering 435
Verifying BGP routing table 436
Verifying BGP peers 436
Clearing BGP neighbors 437
Configuring BGP over an IPSec tunnel interface 438
BGP peering using IPSec tunnel interface IP 438
Best practices and recommendations 440
Use cases 440
Multi-homing to the same ISP 440
BGP AS-Path prepending and AS-Path filters 442
Debugging and troubleshooting 444
FAQs 445
BGP commands 451
address-family 451
aggregate-address 452
bgp add-paths select best 2 454
bgp always-compare-med 455
bgp asnotation dotted 456
bgp asnotation dotted-plus 456
bgp bestpath as-path ignore 457
bgp bestpath as-path multipath-relax 458
bgp bestpath bgp-std-selection 459
bgp bestpath compare-routerid 460
bgp bestpath med confed 461
bgp bestpath med missing-as-worst 462
bgp cluster id 463
bgp confederation 464
bgp confederation peers 465
bgp dampening 466
bgp default local-preference 467
bgp deterministic-med 468
bgp fast-external-fallover 469
bgp graceful-restart restart-time 470
bgp graceful-restart stalepath-time 471
bgp log-neighbor-changes 472
bgp maxas-limit 473
bgp router-id 474
bgp update-group 475
clear bgp 476
disable|enable 478
distance bgp 479
maximum-paths 480
neighbor activate 482
neighbor advertisement-interval 484
neighbor add-paths 485
neighbor add-paths advertise-best 486
neighbor allowas-in 487
neighbor ao 489
neighbor capability orf prefix-list 491
neighbor default-originate 492
neighbor ebgp-multihop 493
neighbor fall-over 495
neighbor fall-over bfd 496

neighbor graceful-shutdown 497
neighbor interface 500
neighbor listen ip-range 501
neighbor local-as 503
neighbor maximum-prefix 505
neighbor next-hop-self 507
neighbor next-hop-unchanged 509
neighbor orf prefix-list in 510
neighbor passive 511
neighbor password 513
neighbor port 514
neighbor remote-as 515
neighbor remove-private-AS 516
neighbor remote-as 518
neighbor replace-rt-asn 519
neighbor route-map 520
neighbor route-reflector-client 522
neighbor send-community 523
neighbor shutdown 525
neighbor slow-peer 527
neighbor soft-reconfiguration inbound 528
neighbor timers 529
neighbor timers connect-retry 530
neighbor ttl-security-hops 531
neighbor update-source 533
neighbor weight 534
network 535
redistribute 536
router bgp 538
show bgp 539
show bgp 543
show bgp community 546
show bgp flap-statistics 549
show bgp ip unicast 551
show bgp neighbor advertised-routes 553
show bgp neighbor paths 555
show bgp neighbor received orf-prefix-list 556
show bgp neighbor received-routes 558
show bgp neighbor routes 559
show bgp neighbors 561
show bgp paths 566
show bgp peer-group summary 567
show bgp summary 569
show bgp l2vpn evpn 572
show bgp l2vpn evpn vni route-type 575
show bgp l2vpn evpn vtep 576
show bgp l2vpn evpn vtep route-type 578
show bgp l2vpn evpn vtep vni 579
show bgp l2vpn evpn vtep vni route-type 580
show bgp unicast protected 582
show ip ospf routes 582
show running-config bgp 586
timers bgp 587
vrf 588

IP Fast Re-Route (IP-FRR) 589

Overview ... 589
Prerequisites ... 589
Limitations ... 590
Fast Traffic Recovery commands ... 590
fast-reroute ... 590
Fast Traffic Recovery commands ... 591
neighbor ... 591
show capacities fast-reroute ... 592
show capacities status fast-reroute ... 592

Route Policies and Route Maps ... 594

Overview ... 594
Route maps ... 594
Match criteria ... 594
Set changes ... 595
IP prefix lists ... 595
AS-path lists for BGP ... 595
Community lists for BGP ... 595
Route flap dampening ... 596
Route redistribution and route maps ... 596
Route policy and route map commands ... 596
General or filtering commands ... 596
ip aspath-list ... 596
ip community-list ... 597
ip prefix-list ... 599
ipv6 prefix-list ... 600
route-map ... 602
continue ... 603
Match commands ... 604
match aspath-list ... 604
match community-list ... 605
match interface ... 606
match ip address prefix-list ... 607
match ip next-hop ... 608
match ip route-source ... 609
match local-preference ... 610
match metric ... 610
match origin ... 611
match route-type ... 612
match source-protocol ... 613
match tag ... 614
match vni ... 615
Set commands ... 616
set as-path exclude ... 616
set as-path prepend ... 617
set community ... 618
set dampening ... 619
set extcommunity ... 620
set ip nexthop ... 621
set ipv6 nexthop global ... 622
set local-broadcast-group ... 623
set local-preference ... 624
set metric ... 625
set origin ... 626
set tag ... 627
set weight ... 628

Show commands 629
show ip aspath-list 629
show ip community-list 630
show ip prefix-list 631
show route-map 632

Equal Cost Multipath (ECMP) 634
Overview 634
Dynamic load balancing (DLB) ECMP Group load balancing modes 634
ECMP commands 635
ip ecmp dynamic-mode 635
show ip ecmp 637

Virtual Router Redundancy Protocol (VRRP) 640
Overview 640
Terminology 640
VRRP operation 641
Multiple VRRP groups 642
VRRP priority 643
VRRP preemption 643
Virtual Router MAC address 644
VRRP and ARP 644
VRRP and MCLAG 644
VRRP tracking 644
High availability 644
VRRP and Neighbor Discovery for IPv6 644
Duplicate address detection (DAD) 645
Guidelines and limitations 645
VRRP commands 645
address 645
authentication 647
bfd <IPV4-ADDR> 649
preempt 650
preempt delay minimum 651
priority 652
router vrrp {enable | disable} 653
no router vrrp 654
show track 654
show track brief 655
show vrrp 656
shutdown 664
timers advertise 665
track (VRRP group) 666
track (VRRP virtual router) 667
track by 668
version 669
vrrp 669
vrrp dual-active-forwarding 670

Inter-Virtual Router Forwarding (IVRF) 672
Static VRF route leaking 672
Dynamic VRF route leaking 673
Procedure to leak routes between VRFs 673
Troubleshooting inter-VRF route leaking 673
IVRF commands 674
address-family 674

ip|ipv6 vrf 675
ipv6 route source interface 678
ip route interface 680
rd 683
redistribute 684
route-target 685
router bgp 687
route-target export 688
router bgp vrf 690
show bgp vpn unicast 691
show bgp info vrf 695
show ip route vrf 696
show ipv6 route vrf 697
vrf 699

Policy Based Routing (PBR) 701

PBR actions 701
PBR policy action and action list 701
PBR action list maximum entries 702
IP versions in an action list 702
Specifying valid next-hop and default-nexthop addresses 703
Hardware path PBR versus software path PBR 703
Hardware versus software path for default-nexthop action 703
Software path and system default route 704
PBR and VRFs 704
PBR, ECMP, and routing protocols 705
PBR, VSX, and VLAN ACLs 705
PBR software path, VSX, and VRRP 705
PBR and next-hop router reachability 705
PBR and VXLAN 706
PBR nexthop over VXLAN L3VNI 706
Ingress policy 707
Egress policy 707
L2VNI vs L3VNI 707
PBR and subinterfaces 708
CLI errors 708
Backup nexthop groups 709
PBR commands 709
apply policy 709
ip nexthop group type backup 710
pbr-action-list 711
pbr-action-list copy 715
pbr-action-list resequence 716
pbr-action-list reset 717
policy 718
primary secondary nexthop 721
show ip nexthop group name type backup 722
show ip nexthop group type backup 722
show pbr 723
show pbr-action-list 725
show running-config current-context 726

IP Directed Broadcast 729

IP Directed Broadcast configuration example 729
IP Directed Broadcast commands 731
copy support-file feature 731

ip directed-broadcast 732
show arp 733
show ip interface 734
show ip directed-broadcast 735

IP Neighbor Flood 737
IP Neighbor Flood commands 737
ip neighbor-flood 737
show ip interface 738
show ip neighbor-flood 739
show running-config 739

Key chain 741
Key chain commands 741
accept-lifetime 741
cryptographic-algorithm 742
key 743
keychain 744
key-string 745
name 746
recv-id 747
send-id 748
send-lifetime 748
show capacities keychain 749
show keychain 750
show running-config keychain 752

IP Client Tracker 754
client track ip 754
client track ip { enable | disable | auto } 755
client track ip client-limit 756
client track ip update-interval 756
client track ip update-method probe 757
show capacities 758
show client ip { count | port | vlan } 759

Routing Information Protocol (RIP) 760
Overview 760
RIPv2 (IPv4) commands 760
Configuration commands 760
router rip 760
Interface commands 761
ip rip 761
ip rip all-ip enable 762
ip rip all-ip disable 763
ip rip all-ip send disable 764
ip rip all-ip receive disable 765
Routing commands 766
enable 766
disable 767
distance 768
maximum-paths 769
redistribute 770
timers update 771
RIPv2 clear commands 772
clear ip rip statistics 772

RIPv2 interface commands ... 773
enable ... 773
disable ... 774
send disable ... 775
receive disable ... 776
RIPv2 show commands ... 777
show capacities rip ... 777
show capacities-status rip ... 778
show ip rip ... 778
show ip rip interface ... 780
show ip rip neighbors ... 781
show ip rip routes ... 782
show ip rip statistics ... 784
show ip rip statistics interface ... 785
show running-config ... 786
RIPng (IPv6) commands ... 787
Configuration commands ... 787
router ripng ... 787
Interface commands ... 788
ipv6 ripng ... 788
Routing commands ... 789
enable ... 789
disable ... 790
distance ... 791
maximum-paths ... 792
redistribute ... 793
timers update ... 794
RIPng clear commands ... 795
clear ipv6 ripng statistics ... 795
RIPng interface commands ... 796
enable ... 796
disable ... 797
send disable ... 798
receive disable ... 799
RIPng show commands ... 800
show capacities ripng ... 800
show capacities-status ripng ... 801
show ipv6 ripng ... 801
show ipv6 ripng interface ... 803
show ipv6 ripng neighbors ... 804
show ipv6 ripng routes ... 805
show ipv6 ripng statistics ... 807
show ipv6 ripng statistics interface ... 808
show running-config ... 809

Support and Other Resources ... 811

Accessing HPE Aruba Networking Support ... 811
Accessing Updates ... 812
Warranty Information ... 812
Regulatory Information ... 812
Documentation Feedback ... 812

Chapter 1
About this document

This document describes features of the AOS-CX network operating system. It is intended for administrators responsible for installing, configuring, and managing HPE Aruba Networking switches on a network.

Applicable products

This document applies to the following products:

- HPE Aruba Networking 6300 Switch Series (JL658A, JL659A, JL660A, JL661A, JL662A, JL663A, JL664A, JL665A, JL666A, JL667A, JL668A, JL762A, R8S89A, R8S90A, R8S91A, R8S92A, S0E91A, S0X44A, S3L75A, S3L76A, S3L77A, S4P41A,S4P42A, S4P43A, S4P44A, S4P45A, S4P46A, S4P47A, S4P48A)
- HPE Aruba Networking 6400 Switch Series (R0X31A, R0X38B, R0X38C, R0X39B, R0X39C, R0X40B, R0X40C, R0X41A, R0X41C, R0X42A, R0X42C, R0X43A, R0X43C, R0X44A, R0X44C, R0X45A, R0X45C, R0X26A, R0X27A, JL741A, S0E48A,S0E48A #0D1, S1T83A, S1T83A #0D1)
- HPE Aruba Networking 8100 Switch Series (R9W94A, R9W95A, R9W96A, R9W97A)
- HPE Aruba Networking 8325 Switch Series (JL624A, JL625A, JL626A, JL627A)
- HPE Aruba Networking 8325H Switch Series (S4B20A, S4B21A, S4B22A, S4B23A, S2T42A, S2T46A, S2T47A, S2T48A)
- HPE Aruba Networking 8325P Switch Series (S0G12A, S4A48A)
- HPE Aruba Networking 8360 Switch Series (JL700A, JL701A, JL702A, JL703A, JL706A, JL707A, JL708A, JL709A, JL710A, JL711A, JL700C, JL701C, JL702C, JL703C, JL706C, JL707C, JL708C, JL709C, JL710C, JL711C, JL704C, JL705C, JL719C, JL718C, JL717C, JL720C, JL722C, JL721C)
- HPE Aruba Networking 9300 Switch Series (R9A29A, R9A30A, R8Z96A, S0F81A, S0F82A, S0F83A)
- HPE Aruba Networking 9300S Switch Series (S0F81A, S0F82A, S0F83A, S0F84A, S0F85A, S0F86A, S0F88A, S0F95A, S0F96A)
- HPE Aruba Networking 10000 Switch Series (R8P13A, R8P14A)

Latest version available online

Updates to this document can occur after initial publication. For the latest versions of product documentation, see the links provided in Support and Other Resources.

Command syntax notation conventions

Convention	Usage
`example-text`	Identifies commands and their options and operands, code examples, filenames, pathnames, and output displayed in a command window. Items that appear like the example text in the previous column are to be entered exactly as shown and are required unless enclosed in brackets (**[ ]**).

Convention	Usage
example-text	In code and screen examples, indicates text entered by a user.
Any of the following: ■ *<example-text>* ■ <example-text> ■ *example-text* ■ *example-text*	Identifies a placeholder—such as a parameter or a variable—that you must substitute with an actual value in a command or in code: ■ For output formats where italic text cannot be displayed, variables are enclosed in angle brackets (**< >**). Substitute the text—including the enclosing angle brackets—with an actual value. ■ For output formats where italic text can be displayed, variables might or might not be enclosed in angle brackets. Substitute the text including the enclosing angle brackets, if any, with an actual value.
|	Vertical bar. A logical **OR** that separates multiple items from which you can choose only one. Any spaces that are on either side of the vertical bar are included for readability and are not a required part of the command syntax.
{ }	Braces. Indicates that at least one of the enclosed items is required.
[ ]	Brackets. Indicates that the enclosed item or items are optional.
... or . . .	Ellipsis: ■ In code and screen examples, a vertical or horizontal ellipsis indicates an omission of information. ■ In syntax using brackets and braces, an ellipsis indicates items that can be repeated. When an item followed by ellipses is enclosed in brackets, zero or more items can be specified.

About the examples

Examples in this document are representative and might not match your particular switch or environment.

The slot and port numbers in this document are for illustration only and might be unavailable on your switch.

Understanding the CLI prompts

When illustrating the prompts in the command line interface (CLI), this document uses the generic term **switch**, instead of the host name of the switch. For example:

```
switch>
```

The CLI prompt indicates the current command context. For example:

```
switch>
```

Indicates the operator command context.

```
switch#
```

Indicates the manager command context.

```
switch(CONTEXT-NAME)#
```

Indicates the configuration context for a feature. For example:

```
switch(config-if)#
```

Identifies the **interface** context.

Variable information in CLI prompts

In certain configuration contexts, the prompt may include variable information. For example, when in the VLAN configuration context, a VLAN number appears in the prompt:

```
switch(config-vlan-100)#
```

When referring to this context, this document uses the syntax:

```
switch(config-vlan-<VLAN-ID>)#
```

Where *<VLAN-ID>* is a variable representing the VLAN number.

Identifying switch ports and interfaces

Physical ports on the switch and their corresponding logical software interfaces are identified using the format: ***member/slot/port*.**

On the HPE Aruba Networking 6300 Switch Series

- *member*: Member number of the switch in a Virtual Switching Framework (VSF) stack. Range: 1 to 10. The primary switch is always member 1. If the switch is not a member of a VSF stack, then member is 1.
- *slot*: Always 1. This is not a modular switch, so there are no slots.
- *port*: Physical number of a port on the switch.

For example, the logical interface **1/1/4** in software is associated with physical port 4 on member 1.

On the HPE Aruba Networking 6400 Switch Series

- *member*: Always 1. VSF is not supported on this switch.
- *slot*: Specifies physical location of a module in the switch chassis.
  - Management modules are on the front of the switch in slots 1/1 and 1/2.
  - Line modules are on the front of the switch starting in slot 1/3.
- *port*: Physical number of a port on a line module.

For example, the logical interface **1/3/4** in software is associated with physical port 4 in slot 3 on member 1.

On the HPE Aruba Networking 8xxx, 93xx, and 10xxx Switch Series

- *member*: Always 1. VSF is not supported on this switch.
- *slot*: Always 1. This is not a modular switch, so there are no slots.
- *port*: Physical number of a port on the switch.

For example, the logical interface **1/1/4** in software is associated with physical port 4 on the switch.

If using breakout cables, the port designation changes to x:y, where x is the physical port and y is the lane when split. For example, the logical interface 1/1/4:2 in software is associated with lane 2 on physical port 4 in slot 1 on member 1.

Identifying modular switch components

- Power supplies are on the front of the switch behind the bezel above the management modules. Power supplies are labeled in software in the format: *member/power supply*:
  - *member*: 1.
  - *power supply*: 1 to 4.
- Fans are on the rear of the switch and are labeled in software as: *member/tray/fan*:
  - *member*: 1.
  - *tray*: 1 to 4.
  - *fan*: 1 to 4.
- Fabric modules are not labeled on the switch but are labeled in software in the format: *member/module*:
  - *member*: 1.
  - *member*: 1 or 2.
- The display module on the rear of the switch is not labeled with a member or slot number.

Chapter 2
Virtual Routing and Forwarding (VRF)

Virtual Routing and Forwarding (VRF) is a Layer 3 level isolation used to achieve Virtual Private Network (VPN). VRF provides overlapping IPs to present and also isolate the routing table from other VPNs in the system.

VRF is a technology that allows multiple instances of a routing table to co-exist within the same router. Because the routing instances are independent, the same or overlapping IP addresses can be used without conflicting with each other. Network functionality is improved because network paths can be segmented without requiring multiple routers.

VRFs at the Core/L3 Gateway device need to be configured on the 10000 Switch series and the Pensando Policy and Services Manager (PSM).

Adding or deleting a VRF

Prerequisites

You must be in the global configuration context, as indicated by the `config` prompt.

Procedure

To configure a VRF, enter the following command.

`vrf` *`<vrf-name>`*

Where <*vrf-name*> is the name of the VRF, up to 32 alphanumeric characters.

The `no` form of the command deletes the VRF and will remove all the configurations from the interfaces which are part of the deleted VRF and move those interfaces to the default VRF.

See the following examples:

```
switch(config)# vrf test
```

```
switch(config)# no vrf test
```

Related topic: vrf (command reference)

IPv4 static route addition or deletion in a VRF

Prerequisites

You must be in the global configuration context, as indicated by the `config` prompt.

Procedure

To add an IPv4 static route to a VRF, enter the following command.

`ip route <dest-ipv4-addr>/<netmask> {<gateway-ip>|<interface>} vrf [<vrf-name>]`

Where

<dest-ipv4-addr>/<netmask>

Specifies the route destination IP address and the network mask length for the destination.

<gateway-ip>|<interface>

Specifies the gateway as either an IP address or an interface.

<vrf-name>

Specifies the VRF name. If no *<vrf-name>* is specified the route is applied to the default VRF.

Use the `no` form of the command to remove an IPv4 static route from the VRF.

See the following example:

```
switch(config)# ip route 20.0.0.0/24 10.0.0.1 vrf test
switch(config)# ip route 20.0.0.0/24 1/1/5 vrf test
```

```
switch(config)# no ip route 20.0.0.0/24 10.0.0.1 vrf test
```

Related topic: ip route vrf (command reference)

IPv6 static route addition or deletion in a VRF

Prerequisites

You must be in the global configuration context, as indicated by the `config` prompt.

Procedure

To add an IPv6 static route to a VRF, enter the following command.

`ipv6 route <dest-ipv6-addr>/<prefix> {<gateway-ip>|<interface>} vrf [<vrf-name>]`

Where

<dest-ipv6-addr>/<prefix>

Specifies the route destination IPv6 address and the network prefix for the destination. For example, `120::/124`.

<gateway-ip>|<interface>

Specifies either the gateway as either an IPv6 address or an interface.

<vrf-name>

Specifies the VRF name. The name can be up to 32 alphanumeric characters. If no *<vrf-name>* is specified the route is applied to the default VRF.

Use the `no` form of the command to delete an IPv6 static route from a VRF.

See the following examples:

```
switch(config)# ipv6 route 120::/124 121::2 vrf test
switch(config)# ipv6 route 120::/124 1/1/9 vrf test
```

```
switch(config)# no ipv6 route 120::/124 121::2 vrf test
```

Related topics: ipv6 route vrf (command reference)

Attaching or detaching a port from a VRF

Prerequisites

You must be in the interface configuration context (`config-if`), or the VLAN interface configuration context (`config-vlan-if`), or the LAG interface configuration context (`config-lag-if`).

Procedure

To attach a port to a VRF, enter the following command in the required context.

`vrf attach <vrf-name>`

Where <*vrf-name*> is the name of the VRF, up to 32 alphanumeric characters.

The `no` form of the command detaches the port from the named VRF and will remove all configurations from the port and attach the port to the default VRF.

See the following examples:

```
switch(config)# interface 1/1/1
switch(config-if)# vrf attach test
```

```
switch(config)# vlan 3
switch(config-vlan)# exit
switch(config)# interface vlan 3
switch(config-if-vlan)# vrf attach test
```

Related topic: vrf attach (command reference)

Client applications running on destination VRF

Client applications such as ping or traceroute must specify a source interface or source IP when they communicates to a destination IP in a destination VRF. If a source interface or source IP is not specified, the client application will chose an internal IP in the range of 240.0.0.0/16 and the server which receives the packet will not be able to communicate back.

Viewing VRF information

Prerequisites

These commands are in the switch context, executed at the `switch#` prompt.

Procedure

To view various aspects of VRF information, use the following commands.

- To view VRF configuration and status information:

`show vrf <VRF-NAME>`

- To view all configured commands including VRF configuration:

`show running-config`

For command details and examples, see the following:

- show vrf
- An example of the VRF information provided by the show running-config command

An example of the VRF information provided by the show running-config command

When a VRF is configured, the output of the **show running-config** command includes information about the VRF configuration.

Example

```
switch# show running-config
Current configuration:
!
!Version AOS-CX 10.0X.XXXX
!
lldp enable
timezone set utc
vrf new
vrf test
!
!
!
aaa authentication login default local
aaa authorization commands default none
!
!
!
!
vlan 1
    no shutdown
interface 1/1/1
    no shutdown
    vrf attach test
    ip address 20.0.0.2/24
    ipv6 address 120::1/124
interface 1/1/2
    no shutdown
    vrf attach new
    ip address 30.0.0.1/24
    ipv6 address 121::1/124
```

VRF commands

ip route vrf

```
ip route <DEST-IPV4-ADDR>/<MASK> [<NEXT-HOP-IP-ADDR>|<INTERFACE>|reject|nullroute]
   vrf <VRF-NAME> [name <WORD>]
no ip route <DEST-IPV4-ADDR>/<MASK> [<NEXT-HOP-IP-ADDR>|<INTERFACE>|reject|nullroute]
   vrf <VRF-NAME> [name <WORD>]
```

Description

Adds the destination IPv4 static route on the specified VRF. If no <*VRF-NAME*> is specified the route is applied to the default VRF.

The **no** form of this command removes the IPv4 static route from the VRF.

Parameter	Description
<DEST-IPV6-ADDR>	Specifies the route destination in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.
<MASK>	Specifies the number of bits in the address mask in CIDR format (**x**), where **x** is a decimal number from 0 to 128.
<NEXT-HOP-IP-ADDR>	Specifies the next hop in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.
<INTERFACE>	Specifies the next hop as an outgoing interface.
`nullroute`	Silently discards packets to the destined route.
`reject`	Discards packets to the destined route and returns an ICMP error to the sender.
`vrf` *<VRF-NAME>*	Specifies a VRF name.
`name` *<WORD>*	Assigns a name to the destination route. Range: 1 to 64 characters.

Examples

```
switch(config)# ip route 20.0.0.0/8 10.20.30.44 vrf myvrf
switch(config)# ip route 20.1.2.0/24 1/1/30 vrf myvrf
switch(config)# ip route 1.2.3.4/32 nullroute vrf myvrf
switch(config)# ip route 1.2.3.4/32 reject vrf myvrf
switch(config)# ip route 20.0.4.0/24 10.0.0.2 vrf test name VRF-Management
```

```
switch(config)# no ip route 20.0.0.0/8 10.20.30.44 vrf myvrf
```

Command History

Release	Modification
10.17	The **name** parameter was introduced.
10.10	Inclusive language update.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H	`config`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325P 8360 9300 9300S 10000		

ipv6 route gc interval

```
ipv6 route-gc-interval <INTERVAL>
no ipv6 route-gc-interval
```

Description

Sets the garbage collection interval timer to remove invalid or old route entries from kernel route cache.

The **no** form of this command resets the garbage collection interval timer to default (30 seconds).

Parameter	Description
<INTERVAL>	Specifies time interval in seconds. Range: 30 to 600. Default: 30.

Examples

Setting garbage collection interval timer to 300:

```
switch(config)# ipv6 route-gc-interval 300
```

Command History

Release	Modification
10.10	Command introduced

Command Information

Platforms	Command context	Authority
All platforms	`config`	Administrators or local user group members with execution rights for this command.

ipv6 route vrf

```
ipv6 route <DEST-IPV6-ADDR>/<PREFIX> [<NEXT-HOP-IP-ADDR>|<INTERFACE>|reject|nullroute]
vrf <VRF-NAME> [name <WORD>]
no ipv6 route <DEST-IPV6-ADDR>/<PREFIX> [<NEXT-HOP-IP-ADDR>|<INTERFACE>|reject|nullroute]
vrf <VRF-NAME> [name <WORD>]
```

Description

Adds an IPv6 static route in the specified VRF. If no *<VRF-NAME>* is specified it is added to the default VRF.

The **no** form of this command removes an IPv6 static route from the VRF.

Parameter	Description
<DEST-IPV6-ADDR>	Specifies an IP address in IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F.
<MASK>	Specifies the number of bits in the address mask in CIDR format (**x**), where **x** is a decimal number from 0 to 128.
<NEXT-HOP-IP-ADDR>	Specifies the next hop in IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F.
<INTERFACE>	Specifies the next hop as an outgoing interface.
`nullroute`	Specifies that packets matching the destination prefix are silently discarded and no ICMP error notification is sent to the sender.
`reject`	Specifies that packets matching the destination prefix are discarded and an ICMP error notification is sent to the sender.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.
`name` *<WORD>*	Assigns a name to the destination route. Range: 1 to 64 characters.

Examples

```
switch(config)# ipv6 route 120::/124 121::2 vrf test
switch(config)# ipv6 route 121::/124 1/1/9 vrf test
switch(config)# ipv6 route 122::/124 nullroute vrf test
switch(config)# ipv6 route 123::/124 reject vrf test
switch(config)# ipv6 route 124::/124 121::3 vrf test name IPv6-VRF-Link
```

```
switch(config)# no ipv6 route 120::/124 121::2 vrf test
```

Command History

Release	Modification
10.17	The **name** parameter was introduced.
10.10	Inclusive language update.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325	`config`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325H 8325P 8360 9300 9300S 10000		

show ip route

`show ip route [<A.B.C.D> | <A.B.C.D/M> | all-vrfs | bgp | connected | local | ospf | static | summary | vrf <VRF-NAME> [vsx-peer] | protected]`

Description

Displays IP route tables.

Parameter	Description
<A.B.C.D>	Display longest prefix match.
<A.B.C.D/M>	Display exact route match.
all-vrfs	Display information for all VRFs.
bgp	Display bgp routes only.
connected	Display connected routes only.
local	Display local routes only.
ospf	Display ospf routes only.
protected	Display protected routes only.
static	Display static routes only.
summary	Display the aggregate count of routes per routing protocol.
vrf *<vrf-name>*	Specify a VRF by VRF name (if no *<VRF-NAME>* is specified, the default VRF is implied.
vsx-peer	

Examples

Showing IP routes:

```
dut1# show ip route all-vrfs

Displaying ip routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF, D - DHCP
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
```

```
              E2 - OSPF external type 2
FRR:          + - Backup Nexthop
              * - data obtainable via 'show bgp ip unicast A.B.C.D/M'

VRF: 1

Prefix        Nexthop  Interface    VRF(egress)    Origin/   Distance/    Age
                                                   Type      Metric
------------------------------------------------------------------------------
--
1.1.1.0/24     -        1/1/22       -              C         [0/0]        -
1.1.1.1/32     -        1/1/22       -              L         [0/0]        -
1.1.2.0/24     -        1/1/23       -              C         [0/0]        -
1.1.2.1/32     -        1/1/23       -              L         [0/0]        -
1.1.5.0/24     -        1/1/25       -              C         [0/0]        -
1.1.5.1/32     -        1/1/25       -              L         [0/0]        -
200.1.0.0/24  1.1.2.2   1/1/23       -              B/E       [20/0]   00h:00m:13s
              1.1.1.2+  1/1/22       -              B/*       [*/0]    00h:00m:13s
200.1.1.0/24  1.1.2.2   1/1/23       -              B/E       [20/0]   00h:00m:13s
              1.1.1.2+  1/1/22       -              B/*       [*/0]    00h:00m:13s
200.1.2.0/24  1.1.2.2   1/1/23       -              B/E       [20/0]   00h:00m:13s
              1.1.1.2+  1/1/22       -              B/*       [*/0]    00h:00m:13s
200.1.3.0/24  1.1.2.2   1/1/23       -              B/E       [20/0]   00h:00m:13s
              1.1.1.2+  1/1/22       -              B/*       [*/0]    00h:00m:13s
200.1.4.0/24  1.1.2.2   1/1/23       -              B/E       [20/0]   00h:00m:13s
              1.1.1.2+  1/1/22       -              B/*       [*/0]    00h:00m:13s
Total Route Count : 11
```

Showing ip protected routes:

```
switch# show ip route protected
all-vrfs  all vrf information
vrf       Specify the VRF name
<cr>
```

Showing protected routes for all vrfs:

Protected routes are only supported on the 9300 Switch Series.

```
switch# show ip route protected all-vrfs

Displaying ip routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF, D - DHCP
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2
FRR:          + - Backup Nexthop
              * - data obtainable via 'show bgp ip unicast A.B.C.D/M'

VRF: default

Prefix           Nexthop     Interface  VRF(egress)  Origin/   Distance/    Age
                                                     Type      Metric
```

```
--------------------------------------------------------------------------------
--
120.120.120.0/24  110.110.110.1  vlan20  -           O        [110/120]
03h:35m:55s
                  130.130.130.1+ vlan22  -           B/*      [*/120] 03h:35m:55s
140.140.140.0/24  130.130.130.1  vlan22  -           O        [110/130]
03h:25m:13s
                  110.110.110.1+ vlan20  -           B/*      [*/130] 03h:25m:13s
211.0.0.0/24      110.110.110.1  vlan20  -           O        [110/122]
03h:25m:13s
                  130.130.130.1+ vlan22  -           B/*      [*/122] 03h:25m:13s

VRF: red

150.150.150.0/24  120.120.120.1  vlan21  -           O        [110/120]
03h:35m:55s
                  140.140.140.1+ vlan23  -           B/*      [*/120] 03h:35m:55s

Total Route Count : 4
```

Showing protected routes for a specific vrf, **vrf red**:

```
switch# show ip route protected vrf red

Displaying ip routes selected for forwarding
Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF, D - DHCP
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2
FRR:          + - Backup Nexthop
              * - data obtainable via 'show bgp ip unicast A.B.C.D/M'

VRF: red

Prefix           Nexthop       Interface VRF(egress)  Origin/   Distance/    Age
                                                      Type      Metric
--------------------------------------------------------------------------------
--
150.150.150.0/24 121.120.120.1  vlan21   -           O         [110/120
03h:35m:55s
                 141.140.140.1+ vlan23   -           B/*       [*/120]   03h:35m:55s

Total Route Count : 1
```

Show IP routes for a specific IP address:

```
dut1# show ip route 200.1.0.0/24 vrf 1

VRF: 1

Prefix           : 200.1.0.0/24                                VRF(egress)
: -
Nexthop          : 1.1.2.2                                     Interface
: 1/1/6
Origin           : bgp                                         Type
: bgp_ext
Distance         : 20                                          Metric
: 0
```

```
Age              : 00h:00m:22s                                    Tag
: 0
Encap Type       : -                                              Encap Details
: -
Nexthop Type     : Primary

Prefix           : 200.1.0.0/24                                   VRF(egress)
: -
Nexthop          : 1.1.1.2                                        Interface
: 1/1/2
Origin           : bgp                                            Type
: *
Distance         : *                                              Metric
: 0
Age              : 00h:00m:22s                                    Tag
: 0
Encap Type       : -                                              Encap Details
: -
Nexthop Type     : Backup

   * - data for backup nexthop obtainable via 'show bgp ip unicast A.B.C.D/M'
```

Showing the ip route summary for all vrfs:

```
core1# show ip route summary all-vrfs

IP Route Table Summary

VRF name :  red
Protocol      Active Routes Protected Routes
------------- ------------- ----------------
connected      3             -
local          3             -
bgp            12            5
```

Command History

Release	Modification
10.14.1000	Added protected paramater to 9300 Switch Series.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ipv6 route

```
show ipv6 route [<X.X.X.X> | <X.X.X.X/M> | all-vrfs | bgp | connected | local | ospf |
static | summary | vrf <vrf-name>] [vsx-peer]
```

Description

Displays IPv6 route tables.

Parameter	Description
<X.X.X.X>	Display exact route match.
<X.X.X.X/M>	Display exact route match.
all-vrfs	Display information for all VRFs.
bgp	Display bgp routes only.
connected	Display connected routes only.
local	Display local routes only.
ospf	Display ospf routes only.
static	Display static routes only.
summary	Display the aggregate count of routes per routing protocol.
vrf <*vrf-name*>	Specify a VRF by VRF name (if no <*VRF-NAME*> is specified, the default VRF is implied.
vsx-peer	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Example

Showing IPv6 route tables:

```
switch# show ipv6 route

Displaying ipv6 routes selected for forwarding

'[x/y]' denotes [distance/metric]

1000::/64, vrf default
        via  vlan2,  [0/0],  connected
1000::1/128, vrf default
        via  vlan2,  [0/0],  local
2000::/64, vrf default
        via  vlan2,  [1/0],  static
2001::/64, vrf default
        via  1000::2,  [1/0],  static
3000:2301::/64, vrf default
        nullroute,  [1/0],  static
4000:2301::/64, vrf default
        reject route,  [1/0],  static
```

Command History

Release	Modification
10.10	Inclusive language update.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show vrf

```
show vrf <VRF-NAME>
show vrf
```

Description

Displays the status and attached interfaces for the specified VRF instance.

The **show vrf** command shows this information for all the VRFs.

Parameter	Description
<VRF-NAME>	Specifies the VRF name. Length: Up to 32 alphanumeric characters.

Examples

Showing VRF information for the test VRF:

```
switch# show vrf test
VRF Configuration:
------------------
VRF Name   : test
        Interfaces              Status
        ------------------------------
        1/1/29                    up
        1/1/30                    up
```

Showing VRF information for all VRFs:

```
switch# show vrf
VRF Configuration:
------------------
VRF Name   : default
        Interfaces              Status
        ----------------------------

VRF Name   : red
        Interfaces              Status
        ----------------------------
        1/1/32                    up

VRF Name   : test
        Interfaces              Status
        ----------------------------
        1/1/29                    up
        1/1/30                    up
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

vrf

```
vrf <VRF-NAME>
no vrf <VRF-NAME>
```

Description

Creates a VRF instance named <*VRF-NAME*> and then enters its context. Use **default** for <*VRF-NAME*> to enter the default VRF configure context.

Except for the default VRF, the **no** form of the command deletes the named VRF instance and any IP configuration for interfaces or SVI linked to default VRF. The default VRF cannot be deleted and a warning is given if attempted. To erase the Route-Distinguisher and Route-Targets, enter the default VRF context and delete them manually one by one.

Only default and mangement VRFs are supported on the following 6300 switches: S3L75A, S3L76A, S3L77A.

Parameter	Description
<VRF-NAME>	Specifies the VRF name. Range: Up to 32 alphanumeric characters. The **mgmt** VRF cannot be used.

Examples

Creating the VRF named **cust_A** and then entering its context:

```
switch(config)# vrf cust_A
```

Entering the **default** VRF context:

```
switch(config)# vrf default
```

Deleting the VRF named **test**:

```
switch(config)# no vrf test
```

Command History

Release	Modification
10.09	Added default VRF information.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config	Administrators or local user group members with execution rights for this command.

vrf attach

```
vrf attach <VRF-NAME>
no vrf attach <VRF-NAME>
```

Description

Attaches the interface to the VRF with the name <*VRF-NAME*>. The command can be entered in several different command contexts.

The **no** form of the command detaches the interface from the named VRF and will remove all configurations from the interface and attach the interface to the default VRF. A warning message is displayed that prompts you whether to proceed: **All Layer 3 configurations associated with the VRF will be deleted. Continue (y/n)?**

Parameter	Description
<VRF-NAME>	Specifies the VRF name. Required. Length: Up to 32 alphanumeric characters.

Examples

```
switch(config)# interface 1/1/29
switch(config-if)# vrf attach test
```

```
switch(config)# vlan 3
switch(config-vlan)# exit
switch(config)# interface vlan 3
switch(config-if-vlan)# vrf attach test
```

```
switch(config)# vrf test
switch(config)# interface lag 3
switch(config-lag-if)# no shutdown
switch(config-lag-if)# vrf attach test
```

```
switch(config)# interface 1/1/29
switch(config-if)# no vrf attach test
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if config-if-vlan config-lag-if config-gre-if	Administrators or local user group members with execution rights for this command.

Chapter 3
Loopback

A loopback interface is a virtual interface supporting IPv4/IPv6 address configuration. The loopback interface can be considered stable because once created, it remains up. The loopback interface can then be configured with an address to use as a reference or identifier independent of the physical interfaces.

- **Device identification**: As long as the router is operational, the state of the loopback interface is always up. Even if only one link to the router is active, the loopback interface can be reached. This functionality makes it possible to identify an active device in the network using the IP address configured on the loopback interface.
- **Routing**: Since the loopback interface is always active, a routing session (such as a BGP session) can continue on an alternate path even if the outbound interface fails. In OSPF, a loopback interface address is advertised as an interface route into the network. This functionality increases reliability by allowing traffic to take alternate paths if there is a link failure. In OSPF and BGP, the router ID can be set to the loopback address to avoid reassignment of the router ID when physical interfaces are added or removed.
- **Device management**: Loopback interface is always reachable and can be used for sending and receiving management information such as logs and SNMP traps without interruption.

Loopback commands

interface loopback

```
interface loopback <INSTANCE>
no interface loopback <INSTANCE>
```

Description

Creates a loopback interface and enters loopback configuration mode.

The **no** form of this command deletes a loopback interface.

Parameter	Description
<INSTANCE>	Selects the loopback interface ID. Range: 0 to 255

Examples

```
switch(config)# interface loopback
switch(config-loopback-if)#
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms except for S3L75A, S3L76A, S3L77A	`config`	

ip address

```
ip address <IPV4-ADDR/MASK> [secondary]
no ip address <IPV4-ADDR/MASK> [secondary]
```

Description

Sets the IPv4 address for a loopback interface.

The **no** form of this command reverses the set of the IPv4 address for a loopback interface.

Parameter	Description
<IPV4-ADDR>	Specifies an IP address in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.
<MASK>	Specifies the number of bits in the address mask in CIDR format (**x**), where **x** is a decimal number from 0 to 128.
secondary	Indicates that the IPv4 address is a secondary address.

Examples

```
switch(config)# interface loopback 1
switch(config-loopback-if)# ip address 16.93.50.2/24
switch(config-loopback-if)# ip address 20.1.1.1/24 secondary
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms except for S3L75A, S3L76A, S3L77A	`config`	

ipv6 address

```
ipv6 address <IPV6-ADDR/MASK>
```

Description

Sets the IPv6 address for a loopback interface.

Parameter	Description
<IPV6-ADDR>	Specifies an IP address in IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F.
<MASK>	Specifies the number of bits in the address mask in CIDR format (**x**), where **x** is a decimal number from 0 to 128.

Examples

switch(config)# **interface loopback 1**

switch(config-loopback-if)# **ipv6 address fd00:5708::f02d:4df6/64**

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms except for S3L75A, S3L76A, S3L77A	`config`	

vrf attach

```
vrf attach <VRF-NAME>
no vrf attach <VRF-NAME>
```

Description

Attaches a non-default VRF to a loopback.

The **no** form of this command deletes a non-default VRF from a loopback and reattaches the default VRF.

Parameter	Description
<VRF-NAME>	Specifies the name of the non-default VRF to be attached/deleted to/from a loopback.

Examples

```
switch(config)# interface loopback 1
switch(config-loopback-if)#vrf attach test
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	

show interface loopback

`show interface loopback [brief | instance <ID>] [vsx-peer]`

Description

This command displays the configuration and status of loopback interfaces.

Parameter	Description
`brief`	Displays brief information about all configured loopback interfaces.
`instance <ID>`	Displays the configuration and status of a loopback interface ID. Range: 1-255
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Examples

```
switch# show interface loopback
 Interface loopback1 is up
 IPv4 address 192.168.1.1/24
 Interface loopback2 is up
 IPv4 address 182.168.1.1/24
```

```
switch# show interface loopback brief
--------------------------------------------------------------------
Loopback      IP Address                                    Status
Interface
```

```
------------------------------------------------------------------------
loopback1   10.1.1.1/24                                          up
loopback1   1111:2222:3333:4444::6666/128                        up
```

```
switch# show interface loopback 1
 Interface loopback1 is up
 IPv4 address 192.168.1.1/24
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms except for S3L75A, S3L76A, S3L77A	Operator (>) or Manager (#)	

Chapter 4
Static routing

Static routes are manually configured. If the network topology is simple enough, you can use static routes for the network's routing requirements. The proper configuration and usage of static routes can improve network performance and ensure bandwidth for important network applications.

The disadvantage of using static routes is that they cannot adapt to network topology changes. If a fault or topological change occurs in the network, the relevant routes will be unreachable and a network administrator must modify the static routes manually.

Static routes can now be assigned descriptive names using the optional name parameter. This feature enhances route management by allowing administrators to identify routes by their purpose rather than just by their network addresses. Names can be up to 64 characters long and help improve readability and maintainability of network configurations.

On the 5420, 6200, 6300, 6400, 8100, and 8360 switch series, a limited number of IPv6 routes with prefixes from 65-127 can be programmed in the ASIC; this allows for hardware/line-rate forwarding for traffic that matches these routes. Any additional IPv6 routes are software forwarded. (Routing performance to destination addresses on these networks may be impacted.) These prefixes are recommended for transit network use only.

Refer to **show capacities** in the ACLs and Classifiers Policy Guide to see the **maximum number of IPv6 routes with prefixes 65-127** that can be configured on the ASIC. Refer to **show capacities-status** in the ACLs and Classifiers Policy Guide to see the maximum number and current consumption.

This limited hardware support for long prefix IPv6 routes applies whether the routes are configured statically or learned dynamically.

These 'long prefix' IPv6 routes are programmed in a switching ASIC's TCAM table instead of the usual route table.

The access switch series are comprised of multiple ASICs of different capabilities with a mixture of line cards with a varying number of capable ASICs. For consistency across platforms, and to support a fair usage scheme amongst TCAM users, there is an imposed maximum long prefix IPv6 route capacity of 510 for all access switch series. Once this hardware limit is exhausted, long prefix IPv6 routes can still be configured/learned but will only be programmed in the software (CPU) path.

Once TCAM/long prefix IPv6 resources are freed up there is no automatic migration of a software-only long prefix IPv6 route to hardware. The route must be removed and recreated, at which point it consumes the freed hardware resource.

Note that since the TCAM table is shared between multiple applications, the 510 long prefix IPv6 route capacity is not guaranteed to be available, depending on current switch configuration/conditions. Use the 'show' resources' command to display current TCAM resource usage & availability.

Each configured/learned long prefix IPv6 route consumes four TCAM entries and these are noted in the 'Long Prefix IPv6 Unicast Route Lookup' row. The maximum 510 TCAM programmed long prefix IPv6 routes would therefore consume 2040 TCAM entries. The remaining Reserved space is necessary for internal TCAM management.

Below is an example of resource usage:

```
show resources
Resource Usage:

Mod  Description
Resource                                                 Used    Reserved     Free
-------------------------------------------------------------------------------
1/3  Long Prefix IPv6 Unicast Route Lookup
          Ingress TCAM Entries                             16        2048
     Total
          Ingress TCAM Entries                             16        2048   18432
          Egress TCAM Entries                               0           0    8192
          Ingress Lookups                                   1                   8
          Ingress Flex Lookups                              0                   1
          Egress Lookups                                    0                   4
          Ingress Policers                                  0                2047
          Egress Policers                                   0                2047
1/4  Long Prefix IPv6 Unicast Route Lookup
          Ingress TCAM Entries                             16        2048
     Total
          Ingress TCAM Entries                             16        2048   18432
          Egress TCAM Entries                               0           0    8192
          Ingress Lookups                                   1                   8
          Ingress Flex Lookups                              0                   1
          Egress Lookups                                    0                   4
          Ingress Policers                                  0                2047
          Egress Policers                                   0                2047
1/5  Long Prefix IPv6 Unicast Route Lookup
          Ingress TCAM Entries                             16        8192
     Total
          Ingress TCAM Entries                             16        8192   57344
          Egress TCAM Entries                               0           0    8192
          Ingress Lookups                                   1                   8
          Ingress Flex Lookups                              0                   1
          Egress Lookups                                    0                   4
          Ingress Policers                                  0                8191
          Egress Policers                                   0                4095
```

Default route

Without a default route, a packet that does not match any routing entries is discarded and an ICMP destination-unreachable packet is sent to the source. A default route is used to forward packets that do not match any routing entry.

The network administrator can configure a default route with the destination as 0.0.0.0 and the mask as 0. The router forwards any packet whose destination address fails to match any entry in the routing table to the next hop of the default static route.

Recursive static routes

A recursive static route is a route for which the next hop is learned from another routing look up (for example, dynamic protocol or from another another static route). For example, the following commands create an unsupported recursive static route:

- `ip route 99.0.0.0/24 30.0.0.2` - Main static route with gateway 30.0.0.2.
- `ip route 30.0.0.0/24 20.0.0.2` - Static route to reach the nexthop of the previously configured route.

Route manager

Route Manager (or Routing Table Manager) is the central repository for all routing information. It maintains the best routes selected by each routing protocol along with various route attributes in the Routing Information Base (RIB). Route Manager also manages routes from various routing protocols and determines the best route to each destination network which needs to be programmed into Forwarding Information Base (FIB).

If the number of routes is larger than 120,000, displaying all routes will require a significant amount of time to display output. A warning message is displayed to suggest using filters to display routes for prefixes for faster results. Use Ctrl+C to interrupt and return to prompt.

Configuration concepts

Before configuring a static route, you must understand the following concepts:

- **Destination address and mask:** In the `ip route` command, an IPv4 address is in dotted-decimal format. A mask can be in the form of mask length - the number of consecutive 1s in the mask.
- **Output interface and next hop address:** When configuring a static route, specify the output interface or next hop address. The next hop address cannot be a local interface IP address or the route configuration will not take effect.
  - Note: On the 6300 and 6400 switch series and the 8100 and 8360 switch series, static routes can be configured with L3 interfaces like route-only-ports, L3 lags, SVIs, hydra-ports and sub-interfaces as nexthop.
- **Other attributes**: You can configure different priorities and administrative distance for different static routes to make route management policies more flexible. For example, specifying the same priority for different routes to the same destination enables load sharing, but specifying different priorities for these routes enables route backup.

Configuration example procedure

Before configuring a static route, complete the following tasks:

- Configure the physical parameters for related interfaces.
- Configure the link-layer attributes for related interfaces.
- Configure the IP addresses for related interfaces.

The number of IPv4 and IPv6 static routes that can be configured is limited to 16K combined.

For the 6300 and 6400 Switch Series, the number of route prefixes can be stored in the routing table, depending on the prefix mask length (data set) being used (for example, /8, /16, /24, etc.). These switches do not support IPv6 addresses with prefixes greater than 64, and for prefixes between 65 and 127 will be software-forwarded.

Basic static route configuration example

The IP addresses and masks of the switches and hosts are displayed here. Static routes are required for interconnection between any two hosts.

Procedure

1. Configure IP addresses for interfaces (details not shown).
2. Configure static routes.
   a. Configure a default route on Switch A.
   ```
   6300# config
   6300(config)#
   <SwitchA> config
   [SwitchA] ip route 0.0.0.0/0 1.1.4.2
   ```
   b. Configure two static routes on Switch B.
   ```
   <SwitchB> config
   [SwitchB] ip route 1.1.2.0/24 1.1.4.1
   [SwitchB] ip route 1.1.3.0/24 1.1.5.6
   ```
   c. Configure a default route on Switch C.
   ```
   6300# config
   6300(config)#
   [SwitchC] ip route 0.0.0.0/0 1.1.5.5
   ```
3. Configure the hosts. The default gateways for hosts A, B, and C are `1.1.2.3`, `1.1.6.1`, and `1.1.3.1`. The configuration procedure is not shown.
4. Display the configuration.
   a. Display the IP routing table of Switch A.
   b. Display the IP routing table of Switch B.
   c. Use the `ping` command on Host B to check the reachability of Host A, assuming Windows XP runs on the two hosts.
   ```
   C:\Documents and Settings\Administrator>ping 1.1.2.2
   Pinging 1.1.2.2 with 32 bytes of data:
   Reply from 1.1.2.2: bytes=32 time=1ms TTL=255
   Reply from 1.1.2.2: bytes=32 time=1ms TTL=255
   Reply from 1.1.2.2: bytes=32 time=1ms TTL=255
   Reply from 1.1.2.2: bytes=32 time=1ms TTL=255
   15
   Ping statistics for 1.1.2.2:
   Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
   Approximate round trip times in milli-seconds:
   Minimum = 1ms, Maximum = 1ms, Average = 1ms
   ```
   d. Use the traceroute command on Host B to check the reachability of Host A.
   ```
   8320# traceroute
   traceroute  Trace the IPv4 route to a device on the network
   traceroute6 Trace the IPv6 route to a device on the network
   8320# traceroute
   A.B.C.D Enter IPv4 address of the device to traceroute
   WORD  Enter host name oo fthe device to traceroute
   8320# traceroute
   ```

Static routing commands

ip route

```
ip route <DEST-IPV4-ADDR>/<NETMASK> [<NEXTHOP-ADDR> | <NEXTHOP-PORT-LAG-VLAN> |
nullroute | reject]
   {distance | name <WORD> | tag <VALUE> | track <WORD> | use-forwarding-address | vrf
   <WORD>]
no ip route <DEST-IPV4-ADDR>/<NETMASK> [<NEXTHOP-ADDR> | <NEXTHOP-PORT-LAG-VLAN> |
nullroute | reject]
```

```
   {distance | name <WORD> | tag <VALUE> | track <WORD> | use-forwarding-address | vrf
   <WORD>]

ip route <DEST-IPV4-ADDR> <NETMASK> [<NEXTHOP-PORT-LAG-VLAN> | nullroute | reject]
   {distance | name <WORD> | tag <VALUE> | vrf <WORD>]
no ip route <DEST-IPV4-ADDR> <NETMASK> [<NEXTHOP-PORT-LAG-VLAN> | nullroute | reject]
   {distance | name <WORD> | tag <VALUE> | vrf <WORD>]
```

Description

Adds an IPv4 static route on the default VRF.

The **no** form of this command deletes a IPv4 static route.

You can configure a maximum of 32 next hops per route.

Parameter	Description
<DEST-IPV4-ADDR>/<NETMASK>	Specifies the IPv4 route destination.
<NETMASK>	Specifies the number of bits in the address mask in CIDR format (**x**), where **x** is a decimal number from 0 to 32.
<NEXTHOP-ADDR>	Specifies the next hop address for reaching the destination in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. Supported on 6300, however S3L75A, S3L76A, and S3L77A do not allow more than one nexthop for a single route.
<NEXTHOP-PORT-LAG-VLAN>	Specifies the next hop as an outgoing interface.
`nullroute`	Specifies that packets matching the destination route are silently discarded and no ICMP error notification is sent to the sender.
`reject`	Specifies that packets matching the destination route are discarded and an ICMP error notification is sent to the sender.
`distance` *<VALUE>*	Specifies the administrative distance to associate with this static route. Range: 1 to 255. Default: 1.
`name` *<WORD>*	Specifies the name for the destination route. Range: 1 to 64 characters.
`tag` *<VALUE>*	Specifies and assigns tag for the route. Range: 1 to 4294967295.
`track` *<WORD>*	Specifies a name for the track object.
`use-forwarding-address`	Specifies the nexthop address as forwarding address.
`vrf` *<WORD>*	Specifies the VRF route leak destination.

Examples

On the HPE Aruba Networking 6400 Switch Series, interface identification differs.

```
switch(config)# ip route 10.0.0.0/24 nullroute
switch(config)# ip route 10.0.1.0/24 reject
switch(config)# ip route 10.0.2.0/24 20.0.0.2
switch(config)# ip route 10.0.3.0/24 1/1/1
```

```
switch(config)# ip route 10.0.3.0/24 1/1/1.110
switch(config)# ip route 10.0.2.0/24 20.0.0.2 distance 4
switch(config)# ip route 10.0.4.0/24 20.0.0.2 name AdminRoute
switch(config)# ip route 10.1.1.5/32 1/1/1 tag 20
switch(config)# ip route 1.1.2.0 255.255.255.0 2.2.2.2 track track1
switch(config)# ip route 10.1.1.6 255.255.255.0 20.1.1.2 use-forwarding-address
switch(config)# ip route 20.0.3.0/24 reject vrf test
switch(config)# ip route 10.0.4.0/24 20.0.0.2 distance 5 name AdminRoute
switch(config)# ip route 1.1.2.0 255.255.255.0 2.2.2.2 track track1 name Tracked-
Critical-Path
switch(config)# ip route 10.1.1.6 255.255.255.0 20.1.1.2 use-forwarding-address
name OSPF-Forward
switch(config)# ip route 10.0.4.0/24 20.0.0.2 tag 15 name TaggedRoute
switch(config)# ip route 20.0.4.0/24 10.0.0.2 vrf test name VRF-Management
```

Command History

Release	Modification
10.17	The **name** parameter was introduced.
10.10	Inclusive language update.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	`config`	Administrators or local user group members with execution rights for this command.

ip route bfd

```
ip route <DEST-IPV4-ADDR>/<NETMASK> [<NEXT-HOP-IP-ADDR> | <INTERFACE> | nullroute |
reject]
   {bfd | distance | name | tag | track | use-forwarding-address | vrf}
no ip route <DEST-IPV4-ADDR>/<NETMASK> [<NEXT-HOP-IP-ADDR> | <INTERFACE> | nullroute |
reject]
   {bfd | distance | name | tag | track | use-forwarding-address | vrf}
```

Description

Enables or disables BFD on the specified static route. To disable BFD, issue the command without the `bfd` option.

Parameter	Description
<DEST-IPV4-ADDR>	Specifies a route destination in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.
<NETMASK>	Specifies the number of bits in the address mask in CIDR format (**x**), where **x** is a decimal number from 0 to 128.
<NEXT-HOP-IP-ADDR>	Specifies the next hop address for reaching the destination in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.

Parameter	Description
<INTERFACE>	Specifies the next hop as an outgoing interface.
bfd	Enables BFD on the static route. Omit this parameter to disable BFD.
name *<WORD>*	Assigns a name to the destination route. Range: 1 to 64 characters.
tag *<1-4294967295>*	Specifies and assigns tag for the route.
vrf *<DST-VRF-NAME>*	The VRF route leak destination.

Examples

On the HPE Aruba Networking 6400 Switch Series, interface identification differs.

Enabling BFD on a static route:

```
switch(config)# interface 1/1/1
switch(config-if)# ip address 20.1.1.2/24
switch(config-if)# no shutdown
switch(config-if)# routing
switch(config-if)# exit
switch(config)# ip route 192.0.0.0/8 20.1.1.1 bfd
switch(config)# ip route 1.1.2.0 255.255.255.0 2.2.2.2 bfd name AdminRoute
```

Disabling BFD on a static route:

```
switch(config)# ip route 192.0.0.0/8 20.1.1.1
```

Command History

Release	Modification
10.17	The **name** parameter was introduced.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300	config	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
9300S 10000		

ip route vrf

```
ip route <DEST-IPV4-ADDR>/<MASK> [<NEXT-HOP-IP-ADDR>|<INTERFACE>|reject|nullroute]
   vrf <VRF-NAME> [name <WORD>]
no ip route <DEST-IPV4-ADDR>/<MASK> [<NEXT-HOP-IP-ADDR>|<INTERFACE>|reject|nullroute]
   vrf <VRF-NAME> [name <WORD>]
```

Description

Adds the destination IPv4 static route on the specified VRF. If no <*VRF-NAME*> is specified the route is applied to the default VRF.

The **no** form of this command removes the IPv4 static route from the VRF.

Parameter	Description
<DEST-IPV6-ADDR>	Specifies the route destination in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.
<MASK>	Specifies the number of bits in the address mask in CIDR format (**x**), where **x** is a decimal number from 0 to 128.
<NEXT-HOP-IP-ADDR>	Specifies the next hop in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.
<INTERFACE>	Specifies the next hop as an outgoing interface.
`nullroute`	Silently discards packets to the destined route.
`reject`	Discards packets to the destined route and returns an ICMP error to the sender.
`vrf <VRF-NAME>`	Specifies a VRF name.
`name <WORD>`	Assigns a name to the destination route. Range: 1 to 64 characters.

Examples

```
switch(config)# ip route 20.0.0.0/8 10.20.30.44 vrf myvrf
switch(config)# ip route 20.1.2.0/24 1/1/30 vrf myvrf
switch(config)# ip route 1.2.3.4/32 nullroute vrf myvrf
switch(config)# ip route 1.2.3.4/32 reject vrf myvrf
switch(config)# ip route 20.0.4.0/24 10.0.0.2 vrf test name VRF-Management
```

```
switch(config)# no ip route 20.0.0.0/8 10.20.30.44 vrf myvrf
```

Command History

Release	Modification
10.17	The **name** parameter was introduced.
10.10	Inclusive language update.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

ipv6 route

```
ipv6 route <DEST-IPV6-ADDR>/<NETMASK> [<NEXTHOP-ADDR> | <NEXTHOP-PORT-LAG-VLAN> |
nullroute | reject]
   {distance | name <WORD> | source <IFNAME> | tag <VALUE> | track <WORD> | vrf <WORD>}
no ipv6 route <DEST-IPV6-ADDR>/<NETMASK> [<NEXTHOP-ADDR> | <NEXTHOP-PORT-LAG-VLAN> |
nullroute | reject]
   {distance | name <WORD> | source <IFNAME> | tag <VALUE> | track <WORD> | vrf <WORD>}
```

Description

Adds an IPv6 static route.

The **no** form of this command deletes an IPv6 static route on the default VRF.

Parameter	Description
<DEST-IPV6-ADDR>	Specifies the route destination in IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F.
<NETMASK>	Specifies the number of bits in the address mask in CIDR format (**x**), where **x** is a decimal number from 0 to 128.
<NEXTHOP-ADDR>	Specifies the next hop in IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F.
<NEXTHOP-PORT-LAG-VLAN>	Specifies the next hop as an outgoing interface.
`nullroute`	Specifies that packets matching the destination route are silently discarded and no ICMP error notification is sent to the sender.
`reject`	Specifies that packets matching the destination route are discarded and an ICMP error notification is sent to the sender.

Parameter	Description
`distance <VALUE>`	Specifies the administrative distance to associate with this static route. Range: 1 to 255. Default: 1.
`name <WORD>`	Specifies the name for the destination route. Range: 1 to 64 characters.
`source <IFNAME>`	Specifies the name for the outgoing interface.
`tag <VALUE>`	Specifies and assigns tag for the route. Range: 1 to 4294967295.
`track <WORD>`	Specifies the object monitor. Range: 1 to 64 characters.
`vrf <WORD>`	Specifies the VRF route leak destination.

Usage

On the HPE Aruba Networking 6200, 6300, 6400, 8100, and 8360 Switch series, a limited number of IPv6 routes with prefixes from 65-127 can be programmed in the ASIC; this allows for hardware/line-rate forwarding for traffic that matches these routes. Any additional IPv6 routes are software forwarded. (Routing performance to destination addresses on these networks may be impacted.) Use the **show capacities l3-resources** command to see the **maximum number of IPv6 routes with these prefix lengths** that can be configured in the ASIC. These prefixes are recommended for transit network use only.

Refer to **show capacities** in the ACLs and Classifier Policies Guide to see the **maximum number of IPv6 routes with prefixes 65-127** that can be configured on the ASIC. Refer to **show capacities-status** in the ACLs and Classifier Policies Guide to see the maximum number and current consumption.

This limited hardware support for long prefix IPv6 routes applies whether the routes are configured statically or learned dynamically.

Examples

On the HPE Aruba Networking 6400 Switch Series, interface identification differs.

```
switch(config)# ipv6 route 120::/124 nullroute
switch(config)# ipv6 route 122::/124 1/1/1
switch(config)# ipv6 route 122::/124 1/1/1.110
switch(config)# ipv6 route 123::/124 120::1
switch(config)# ipv6 route 122::/124 nullroute distance 5
switch(config)# ipv6 route 3001::1/128 nullroute tag 10
switch(config)# ipv6 route 122::/124 reject distance 5
switch(config)# ipv6 route 3001::1/128 reject tag 10
switch(config)# ipv6 route 123::/124 120::1 distance 6
switch(config)# ipv6 route 124::/124 120::1 name IPv6ToHost
switch(config)# ipv6 route 3002::1/128 1000::2 tag 20
switch(config)# ipv6 route 2000::/64 1000::2 track track1
switch(config)# ipv6 route 123::/124 reject vrf test
switch(config)# ipv6 route 123::/124 nullroute distance 6 name IPv6-Block
switch(config)# ipv6 route 3002::1/128 nullroute tag 11 name IPv6-Block
switch(config)# ipv6 route 123::/124 reject distance 6 name IPv6-Reject-Route
switch(config)# ipv6 route 3002::1/128 reject tag 11 name IPv6-ICMP-Error
switch(config)# ipv6 route 125::/124 120::1 distance 7 name V6ToServer
switch(config)# ipv6 route 3004::1/128 1000::3 tag 40 name Backup-Link
switch(config)# ipv6 route 2001::/64 1000::2 track track1 name IPv6-Monitored-Link
```

```
switch(config)# ipv6 route 124::/124 121::3 vrf test name IPv6-VRF-Link
```

Command History

Release	Modification
10.17	The **name** parameter was introduced.
10.10	Inclusive language update.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	`config`	Administrators or local user group members with execution rights for this command.

show ip rib

`show ip rib <FILTER> [vrf <VRF-NAME>]`

Description

Shows the IPv4 Routing Information Base (RIB) of VRF with name (**<VRF-NAME>**). If VRF name is not specified, default VRF routes are displayed.

Parameter	Description
<FILTER>	Selects filter, see Usage section.
`vrf <VRF-NAME>`	Specifies the VRF name.

Usage

There are sub-options available within this command:

- **A.B.C.D:** Shows longest prefix match.
- **A.B.C.D/M:** Shows exact route match.
- **all-vrfs:** Shows all VRF information.
- **bgp:** Shows BGP routes only.
- **connected:** Shows connected routes only.
- **host-routes:** Shows ARP/ND learned routes only.
- **connected:** Shows connected routes only.
- **local:** Shows local routes only.
- **ospf:** Shows OSPF routes only.
- **rip:** Shows RIP routes only.
- **static:** Shows static routes only.
- **summary:** Shows aggregate count of routes per routing protocol.

- **vrf:** Specifies the VRF name.
- **selected:** Shows routes selected for forwarding only.
- **non-selected:** Shows routes not selected for forwarding only.

The output of the **show ip rib** commands are not available in information generated by the **show tech files** command. This information is now available in the file ipv4_rib_dump.gz, which can be generated using the command **sudo ovs-appctl -t hpe-routing hpe-metaswitch/show_ip_rib ipv4 all-vrfs**.

Examples

Showing IPv4 routes in RIB for HPE Aruba Networking 8325, 10000, 8360, and 9300/9300S Switch series:

```
switch# show ip rib
Origin Codes: R - RIP, O - OSPFv2, B - BGP
              C - connected, S - static, H - host-routes
Type Codes:   E - External BGP, I - Internal BGP, IA - OSPF inter area
              E1 - OSPF external type 1, E2 - OSPF external type 2
* indicates selected for forwarding

VRF: default

Prefix          Nexthop      Interface   VRF       Origin/ Distance/ Age
                                                       Type    Metric
-------------------------------------------------------------------------------
--
*1.1.1.1/32     -            1/1/33      -         H      [1/0]
00h:00m:07s
*10.0.0.0/30    -            1/1/1       -         S      [20/0]
0d:10h:01m:41s
*10.0.1.0/30    -            1/1/1       -         B/I    [200/0]
2d:20h:01m:42s
*10.1.64.0/18   -            loopback2   -         C      [0/0]      -
*10.2.64.0/18   10.0.0.3     lag1        -         O/E1   [110/25]
1d:05h:03m:43s
*10.2.64.0/18   20.10.0.1    vlan100     -         O/E1   [110/25]
0d:05h:03m:43s
*20.1.2.3/32    2.2.2.2      1/1/4       vrf_red   B/E    [20/0]
2d:10h:01m:45s
*30.1.3.0/24    -            reject      -         S      [1/0]
33d:10h:01m:43s
*50.10.13.0/24  -            reject      -         S      [1/0]
12d:10h:01m:44s
*61.1.1.2/32    4.4.4.4      1/1/5       -         B/I    [200/0]
1d:11h:01m:45s
*62.1.1.3/32    5.5.5.5      1/1/6       -         B/I    [200/0]
0d:12h:01m:45s
*193.0.0.2/32   50.0.0.2     1/1/2       -         S      [1/0]
0d:04h:01m:43s
 193.0.0.2/32   56.0.0.3     1/1/3       -         O/E1   [110/25]
0d:04h:03m:43s

Total Route Count : 14
```

Showing IPv4 routes in RIB for all other switches:

```
switch# show ip rib
Origin Codes: R - RIP, O - OSPFv2, B - BGP
              C - connected, S - static, D - DHCP
Type Codes:   E - External BGP, I - Internal BGP, IA - OSPF inter area
```

```
              E1 - OSPF external type 1, E2 - OSPF external type 2
* indicates selected for forwarding

VRF: default

Prefix          Nexthop      Interface   VRF       Origin/ Distance/ Age
                                                       Type    Metric
-------------------------------------------------------------------------------
--
*1.1.1.1/32     -            1/1/33      -         H      [1/0]
00h:00m:07s
*10.0.0.0/30    -            1/1/1       -         S      [20/0]
0d:10h:01m:41s
*10.0.1.0/30    -            1/1/1       -         B/I    [200/0]
2d:20h:01m:42s
*10.1.64.0/18   -            loopback2   -         C      [0/0]      -
*10.2.64.0/18   10.0.0.3     lag1        -         O/E1   [110/25]
1d:05h:03m:43s
*10.2.64.0/18   20.10.0.1    vlan100     -         O/E1   [110/25]
0d:05h:03m:43s
*20.1.2.3/32    2.2.2.2      1/1/4       vrf_red   B/E    [20/0]
2d:10h:01m:45s
*30.1.3.0/24    -            reject      -         S      [1/0]
33d:10h:01m:43s
*50.10.13.0/24  -            reject      -         S      [1/0]
12d:10h:01m:44s
*61.1.1.2/32    4.4.4.4      1/1/5       -         B/I    [200/0]
1d:11h:01m:45s
*62.1.1.3/32    5.5.5.5      1/1/6       -         B/I    [200/0]
0d:12h:01m:45s
*193.0.0.2/32   50.0.0.2     1/1/2       -         S      [1/0]
0d:04h:01m:43s
 193.0.0.2/32   56.0.0.3     1/1/3       -         O/E1   [110/25]
0d:04h:03m:43s

Total Route Count : 14
```

Showing IPv4 exact route match in RIB:

```
switch# show ip rib 10.0.0.0/30

VRF : default

Prefix          : 10.0.0.0/30           VRF(egress)        :  -
Nexthop         :  -                    Interface          :  1/1/1
Origin          : Connected             Type               :  -
Distance        : 0                     Metric             :  0
Age             : -                     Tag                :  0
Selected        : Yes                   Recursive Nexthop  :  No
```

Showing IPv4 RIB summary:

```
switch# show ip rib summary

IPv4 RIB Table Summary

VRF name :      default
  Protocol      RIB Routes
--------------- -------------
```

```
connected       1010
local           1011
host-routes     2
static          4
ospfv2          509
bgp             9014
selected        10008
non-selected    1518
```

Command History

Release	Modification
10.12	Added **host-routes** parameter and updated output.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

show ipv6 rib

`show ipv6 rib <FILTER> [vrf <VRF-NAME>]`

Description

Shows the IPv6 Routing Information Base (RIB) of VRF with name (**<VRF-NAME>**). If VRF name is not specified, default VRF routes are displayed.

Parameter	Description
<FILTER>	Selects filter, see usage section.
`vrf <VRF-NAME>`	Shows routes in the VRF and specifies VRF name.

Usage

There are sub-options available within this command:

- **X:X: :X:X:** Shows longest prefix match.
- **X:X: :X:X/M:** Shows exact route match.
- **all-vrfs:** Shows all VRF information.
- **bgp:** Shows BGP routes only.
- **connected:** Shows connected routes only.
- **host-routes:** Show ARP/ND learned routes only.
- **local:** Shows local routes only.
- **ospf:** Shows OSPF routes only.

- **rip:** Shows RIP routes only.
- **static:** Shows static routes only.
- **summary:** Shows aggregate count of routes per routing protocol.
- **vrf:** Specifies the VRF name.
- **selected:** Shows routes selected for forwarding only.
- **non-selected:** Shows routes not selected for forwarding only.

Examples

Showing IPv6 routes in RIB for HPE Aruba Networking 8325, 10000, 8360, and 9300/9300S Switch series:

```
switch# show ipv6 rib
Origin Codes: R - RIPng, O - OSPFv3, B - BGP
              C - connected, S - static, H - host-routes
Type Codes:   E - External BGP, I - Internal BGP, IA - OSPF inter area
              E1 - OSPF external type 1, E2 - OSPF external type 2
* indicates selected for forwarding

VRF: default

Prefix               Nexthop     Interface  VRF    Origin/ Distance/ Age
                                                       Type    Metric
------------------------------------------------------------------------------
--
*1::2/128            1::2        1/1/1      -      H      [1/0]        00h:00m:06s
*1000::/64           -           1/1/1      -      C      [0/0]     -
*1000::8/128         -           1/1/1      -      L      [0/0]     -
*1001:db8::/32       1000::10    1/1/1      -      B/I    [200/0]   1d:20h:01m:42s
*2000::/64           fe80::3182  vlan100    -      S      [1/0]     2d:05h:03m:43s
 2000::/64           fe80::1241  1/1/1      -      O/E1   [110/25]  0d:05h:03m:43s
*2000::2000:0:0:0/67 fe80::1111  lag1       Green  B/E    [20/0]    1d:10h:01m:45s
*3001::0/64          -           vlan100    -      C      [0/0]     -
*3001::1/128         -           vlan100    -      L      [0/0]     -
*6101::0/64          -           nullroute  -      S      [1/0]
12d:10h:01m:43s

Total Route Count : 10
```

Showing IPv6 routes in RIB for the other switch series:

```
switch# show ipv6 rib
Origin Codes: R - RIPng, O - OSPFv3, B - BGP
              C - connected, S - static, D- DHCP
Type Codes:   E - External BGP, I - Internal BGP, IA - OSPF inter area
              E1 - OSPF external type 1, E2 - OSPF external type 2
* indicates selected for forwarding

VRF: default

Prefix               Nexthop     Interface  VRF    Origin/ Distance/ Age
                                                       Type    Metric
------------------------------------------------------------------------------
--
*1000::/64           -           1/1/1      -      C      [0/0]     -
*1000::8/128         -           1/1/1      -      L      [0/0]     -
*1001:db8::/32       1000::10    1/1/1      -      B/I    [200/0]   1d:20h:01m:42s
*2000::/64           fe80::3182  vlan100    -      S      [1/0]     2d:05h:03m:43s
```

```
 2000::/64           fe80::1241  1/1/1      -      O/E1   [110/25]  0d:05h:03m:43s
*2000::2000:0:0:0/67 fe80::1111  lag1       Green  B/E    [20/0]    1d:10h:01m:45s
*3001::0/64          -           vlan100    -      C      [0/0]     -
*3001::1/128         -           vlan100    -      L      [0/0]     -
*6101::0/64          -           nullroute  -      S      [1/0]
12d:10h:01m:43s

Total Route Count : 10
```

Showing IPv6 exact route match in RIB:

```
switch# show ipv6 rib 2000::2000:0:0:0

VRF : default

Prefix           : 2000::2000:0:0:0/67         VRF(egress)        :  Green
Nexthop          : fe80::1111                  Interface          :  lag1
Origin           : BGP                         Type               :  External
Distance         : 20                          Metric             :  0
Age              : 1d:10h:01m:45s              Tag                :  20
Selected         : Yes                         Recursive Nexthop  :  Yes
```

Showing IPv6 RIB summary:

```
switch# show ipv6 rib summary

IPv6 RIB Table Summary

VRF name :  default
  Protocol      RIB Routes
-------------- -------------
connected      1009
local          1010
static         3
host-routes    1
ospfv3         508
bgp            1013

selected       10004
non-selected   1527
```

Command History

Release	Modification
10.12	Added **host-routes** parameter and updated output.
10.10	Inclusive language update.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

ipv6 route vrf

```
ipv6 route <DEST-IPV6-ADDR>/<PREFIX> [<NEXT-HOP-IP-ADDR>|<INTERFACE>|reject|nullroute]
vrf <VRF-NAME> [name <WORD>]
no ipv6 route <DEST-IPV6-ADDR>/<PREFIX> [<NEXT-HOP-IP-ADDR>|<INTERFACE>|reject|nullroute]
vrf <VRF-NAME> [name <WORD>]
```

Description

Adds an IPv6 static route in the specified VRF. If no <*VRF-NAME*> is specified it is added to the default VRF.

The **no** form of this command removes an IPv6 static route from the VRF.

Parameter	Description
<DEST-IPV6-ADDR>	Specifies an IP address in IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F.
<MASK>	Specifies the number of bits in the address mask in CIDR format (**x**), where **x** is a decimal number from 0 to 128.
<NEXT-HOP-IP-ADDR>	Specifies the next hop in IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F.
<INTERFACE>	Specifies the next hop as an outgoing interface.
`nullroute`	Specifies that packets matching the destination prefix are silently discarded and no ICMP error notification is sent to the sender.
`reject`	Specifies that packets matching the destination prefix are discarded and an ICMP error notification is sent to the sender.
`vrf <VRF-NAME>`	Specifies the name of a VRF. Default: default.
`name <WORD>`	Assigns a name to the destination route. Range: 1 to 64 characters.

Examples

```
switch(config)# ipv6 route 120::/124 121::2 vrf test
switch(config)# ipv6 route 121::/124 1/1/9 vrf test
switch(config)# ipv6 route 122::/124 nullroute vrf test
switch(config)# ipv6 route 123::/124 reject vrf test
switch(config)# ipv6 route 124::/124 121::3 vrf test name IPv6-VRF-Link
```

```
switch(config)# no ipv6 route 120::/124 121::2 vrf test
```

Command History

Release	Modification
10.17	The **name** parameter was introduced.
10.10	Inclusive language update.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

show ip route

`show ip route [<`**`A.B.C.D`**`> | <`**`A.B.C.D/M`**`> | all-vrfs | bgp | connected | local | ospf | static | summary | vrf <`*`VRF-NAME`*`> [vsx-peer] | protected]`

Description

Displays IP route tables.

Parameter	Description
`<A.B.C.D>`	Display longest prefix match.
`<A.B.C.D/M>`	Display exact route match.
`all-vrfs`	Display information for all VRFs.
`bgp`	Display bgp routes only.
`connected`	Display connected routes only.
`local`	Display local routes only.
`ospf`	Display ospf routes only.
`protected`	Display protected routes only.
`static`	Display static routes only.
`summary`	Display the aggregate count of routes per routing protocol.
`vrf <vrf-name>`	Specify a VRF by VRF name (if no <*VRF-NAME*> is specified, the default VRF is implied.

Parameter	Description
vsx-peer	

Examples

Showing IP routes:

```
dut1# show ip route all-vrfs

Displaying ip routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF, D - DHCP
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2
FRR:          + - Backup Nexthop
              * - data obtainable via 'show bgp ip unicast A.B.C.D/M'

VRF: 1

Prefix        Nexthop  Interface    VRF(egress)    Origin/   Distance/    Age
                                                   Type      Metric
-------------------------------------------------------------------------------
--
1.1.1.0/24     -        1/1/22       -              C         [0/0]        -
1.1.1.1/32     -        1/1/22       -              L         [0/0]        -
1.1.2.0/24     -        1/1/23       -              C         [0/0]        -
1.1.2.1/32     -        1/1/23       -              L         [0/0]        -
1.1.5.0/24     -        1/1/25       -              C         [0/0]        -
1.1.5.1/32     -        1/1/25       -              L         [0/0]        -
200.1.0.0/24  1.1.2.2   1/1/23       -              B/E       [20/0]   00h:00m:13s
              1.1.1.2+  1/1/22       -              B/*       [*/0]    00h:00m:13s
200.1.1.0/24  1.1.2.2   1/1/23       -              B/E       [20/0]   00h:00m:13s
              1.1.1.2+  1/1/22       -              B/*       [*/0]    00h:00m:13s
200.1.2.0/24  1.1.2.2   1/1/23       -              B/E       [20/0]   00h:00m:13s
              1.1.1.2+  1/1/22       -              B/*       [*/0]    00h:00m:13s
200.1.3.0/24  1.1.2.2   1/1/23       -              B/E       [20/0]   00h:00m:13s
              1.1.1.2+  1/1/22       -              B/*       [*/0]    00h:00m:13s
200.1.4.0/24  1.1.2.2   1/1/23       -              B/E       [20/0]   00h:00m:13s
              1.1.1.2+  1/1/22       -              B/*       [*/0]    00h:00m:13s
Total Route Count : 11
```

Showing ip protected routes:

```
switch# show ip route protected
all-vrfs  all vrf information
vrf       Specify the VRF name
<cr>
```

Showing protected routes for all vrfs:

Protected routes are only supported on the 9300 Switch Series.

```
switch# show ip route protected all-vrfs

Displaying ip routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF, D - DHCP
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2
FRR:          + - Backup Nexthop
              * - data obtainable via 'show bgp ip unicast A.B.C.D/M'

VRF: default

Prefix            Nexthop     Interface  VRF(egress)  Origin/   Distance/    Age
                                                       Type      Metric
-------------------------------------------------------------------------------
--
120.120.120.0/24  110.110.110.1  vlan20  -            O        [110/120]
03h:35m:55s
                  130.130.130.1+ vlan22  -            B/*      [*/120] 03h:35m:55s
140.140.140.0/24  130.130.130.1  vlan22  -            O        [110/130]
03h:25m:13s
                  110.110.110.1+ vlan20  -            B/*      [*/130] 03h:25m:13s
211.0.0.0/24      110.110.110.1  vlan20  -            O        [110/122]
03h:25m:13s
                  130.130.130.1+ vlan22  -            B/*      [*/122] 03h:25m:13s

VRF: red

150.150.150.0/24  120.120.120.1  vlan21  -            O        [110/120]
03h:35m:55s
                  140.140.140.1+ vlan23  -            B/*      [*/120] 03h:35m:55s

Total Route Count : 4
```

Showing protected routes for a specific vrf, **vrf red**:

```
switch# show ip route protected vrf red

Displaying ip routes selected for forwarding
Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF, D - DHCP
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2
FRR:          + - Backup Nexthop
              * - data obtainable via 'show bgp ip unicast A.B.C.D/M'

VRF: red

Prefix           Nexthop       Interface VRF(egress)  Origin/   Distance/    Age
                                                      Type      Metric
-------------------------------------------------------------------------------
--
150.150.150.0/24 121.120.120.1  vlan21   -            O         [110/120
03h:35m:55s
                 141.140.140.1+ vlan23   -            B/*       [*/120]   03h:35m:55s

Total Route Count : 1
```

Show IP routes for a specific IP address:

```
dut1# show ip route 200.1.0.0/24 vrf 1

VRF: 1

Prefix           : 200.1.0.0/24                                  VRF(egress)
: -
Nexthop          : 1.1.2.2                                       Interface
: 1/1/6
Origin           : bgp                                           Type
: bgp_ext
Distance         : 20                                            Metric
: 0
Age              : 00h:00m:22s                                   Tag
: 0
Encap Type       : -                                             Encap Details
: -
Nexthop Type     : Primary

Prefix           : 200.1.0.0/24                                  VRF(egress)
: -
Nexthop          : 1.1.1.2                                       Interface
: 1/1/2
Origin           : bgp                                           Type
: *
Distance         : *                                             Metric
: 0
Age              : 00h:00m:22s                                   Tag
: 0
Encap Type       : -                                             Encap Details
: -
Nexthop Type     : Backup

   * - data for backup nexthop obtainable via 'show bgp ip unicast A.B.C.D/M'
```

Showing the ip route summary for all vrfs:

```
core1# show ip route summary all-vrfs

IP Route Table Summary

VRF name :  red
Protocol      Active Routes Protected Routes
------------- ------------- ----------------
connected      3             -
local          3             -
bgp            12            5
```

Command History

Release	Modification
10.14.1000	Added protected paramater to 9300 Switch Series.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ip rib host-routes

`show ip rib host-routes[vrf <VRF-NAME>]`

Description

Shows the IPv4 host-routes Routing Information Base (RIB) routes of VRF with name (`<VRF-NAME>`). If VRF name is not specified, default VRF routes are displayed.

Parameter	Description
<VRF>	Displays the routes in the VRF.
`vrf <VRF-NAME>`	Specifies the VRF name.

Examples

Showing IPv4 routes in RIB:

```
switch# show ipv4 rib host-routes
Displaying ipv4 routes in RIB
Origin Codes: R - RIP, O - OSPFv2, B - BGP
              C - connected, S - static, L - local, H - host-routes
Type Codes:   E - External BGP, I - Internal BGP
              IA - OSPF inter area, ia - OSPF intra area
              E1 - OSPF external type 1, E2 - OSPF external type 2
* indicates selected for forwarding
VRF: default
Prefix              Nexthop          Interface     VRF(egress)       Origin/
Distance/    Age
                                                                     Type
Metric
-------------------------------------------------------------------------------
-----------------------
*1.1.1.2/32         1.1.1.2          1/1/1         -               H         [1/0]
       00h:08m:35s
Total Route Count : 1
```

Command History

Release	Modification
10.12	Command introduced.

Command Information

Platforms	Command context	Authority
9300 10000 8325 8360	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

show ipv6 route

`show ipv6 route [<X.X.X.X> | <X.X.X.X/M> | all-vrfs | bgp | connected | local | ospf | static | summary | vrf <vrf-name>] [vsx-peer]`

Description

Displays IPv6 route tables.

Parameter	Description
<X.X.X.X>	Display exact route match.
<X.X.X.X/M>	Display exact route match.
`all-vrfs`	Display information for all VRFs.
`bgp`	Display bgp routes only.
`connected`	Display connected routes only.
`local`	Display local routes only.
`ospf`	Display ospf routes only.
`static`	Display static routes only.
`summary`	Display the aggregate count of routes per routing protocol.
`vrf <vrf-name>`	Specify a VRF by VRF name (if no <*VRF-NAME*> is specified, the default VRF is implied.
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Example

Showing IPv6 route tables:

```
switch# show ipv6 route

Displaying ipv6 routes selected for forwarding

'[x/y]' denotes [distance/metric]
```

```
1000::/64, vrf default
        via  vlan2,  [0/0],  connected
1000::1/128, vrf default
        via  vlan2,  [0/0],  local
2000::/64, vrf default
        via  vlan2,  [1/0],  static
2001::/64, vrf default
        via  1000::2,  [1/0],  static
3000:2301::/64, vrf default
        nullroute,  [1/0],  static
4000:2301::/64, vrf default
        reject route,  [1/0],  static
```

Command History

Release	Modification
10.10	Inclusive language update.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ipv6 rib host-routes

`show ipv6 rib host-routes[vrf <VRF-NAME>]`

Description

Shows the IPv6 host-routes Routing Information Base (RIB) routes of VRF with name (**<VRF-NAME>**). If VRF name is not specified, default VRF routes are displayed.

Parameter	Description
<VRF>	Displays the routes in the VRF.
`vrf` *<VRF-NAME>*	Specifies the VRF name.

Examples

Showing IPv6 routes in RIB:

```
switch# show ipv6 rib host-routes
Displaying ipv6 routes in RIB
Origin Codes: R - RIP, O - OSPFv2, B - BGP
              C - connected, S - static, L - local, H - host-routes
Type Codes:   E - External BGP, I - Internal BGP
              IA - OSPF inter area, ia - OSPF intra area
              E1 - OSPF external type 1, E2 - OSPF external type 2
* indicates selected for forwarding
VRF: default
Prefix              Nexthop          Interface     VRF(egress)       Origin/
Distance/    Age
                                                                     Type
Metric
------------------------------------------------------------------------------
-----------------------
*1::2/128           1::2             1/1/1         -                 H
[1/0]      00h:48m:08s

Total Route Count : 1
```

Command History

Release	Modification
10.12	Command introduced.

Command Information

Platforms	Command context	Authority
9300 10000 8325 8360	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

Chapter 5

Open Shortest Path First version 2

Open Shortest Path First version 2 (OSPFv2) is a routing protocol described in RFC2328. It is a Link State-based IGP (Interior Gateway Protocol) routing protocol applied to routers grouped into OSPF areas identified by the routing configuration on each routing switch. It is widely used in medium to large-sized enterprise networks.

Overview

The characteristics of OSPFv2 are:

- Provides a loop-free topology using SPF algorithm.
- Allows hierarchical routing using area 0 (backbone area) as the top of the hierarchy.
- Supports load balancing with equal cost routes for the same destination.
- OSPFv2 is a classless protocol and allows for a hierarchical design with VLSM (Variable Length Subnet Masking) and route summarization.
- Provides authentication of routing messages.
- Scales easily using the concept of OSPF areas.
- Provides fast convergence with triggered, incremental updates via LSAs.

Some OSPFv2 configuration is done in the global configuration context, others in the router ospf context, or in the interface configuration context, or in the vlink context. OSPFv2 can be configured on L3 ports, VLAN interfaces, LAG interfaces, and loopback interfaces. All such configurations work in the mentioned interfaces context. OSPFv2 mandates the associated interface to be a routed interface.

Supported features

- OSPFv2 neighbor adjacency, Hello protocol, multiple areas, Inter-area routing
- AS external routes, stub areas, totally stubby areas, NSSA, ABR, ASBR
- Designated Router/Backup Designated Router
- Point-to-point interfaces/broadcast interfaces
- Virtual Links
- Bidirectional Forwarding Detection (BFD) - refer to the *High Availability Guide* for additional information
- Equal-cost multipath
- Null authentication, Simple password authentication, MD5 authentication
- Area range aggregation - Type-3/Type-7
- External route aggregation - Type-5 LSAs
- Graceful Restart - un-planned, restart interval, Graceful Restart Helper
- Stub router advertisement
- SPF throttling
- LSA Throttling
- SPF throttling

- Configuration of interface parameters such as priority, cost, hello-interval, dead-interval, retransmit-interval, transit-delay, etc.
- Configuration of virtual link parameters such as hello-interval, dead-interval, retransmit-interval, transit-delay, etc.
- Route redistribution
- Passive interfaces
- Multi VRF support with each VRF having up to eight OSPF process instances
- Congestion control (prioritizing hello packets, inactivity timer reset, and adjacency throttling)
- IP unnumbered interfaces

Listed below are a few recommendations while configuring IP unnumbered interfaces:

- Enable OSPF on the lending interface in the same area as the borrowing interface.
- If different unnumbered interfaces are used for different areas, then configure them to use different loop backs.
- Configure the address on the lending loopback as /32.
- OSPFv3 can also be enabled on an IP unnumbered interface. OSPFv3 neighborship can be established on IP unnumbered interfaces as it works over IPv6 link-local address.

How OSPFv2 protocol works

The protocol uses Link State Advertisements (LSAs) transmitted by each router to update neighboring routers regarding its interfaces and the routes available through those interfaces. Each routing switch in an area also maintains a link-state database (LSDB) that describes the area topology. (All routers in a given OSPF area have identical LSDBs.) The routing switches used to connect areas to each other flood summary link LSAs and external link LSAs to neighboring OSPF areas to update them regarding available routes. Through this means, each OSPF router determines the shortest path between itself and a desired destination router in the same OSPF domain (Autonomous System (AS)).

OSPFv2 concepts

The following sections describe OSPFv2 Link-State types, router types, and area types.

OSPFv2 Link-state advertisement (LSA) types

OSPFv2 sends routing information in Link-State Advertisements (LSAs). The switches support the following types of LSAs.

- **Router LSA—Type-1 LSA.** Describes the states of the router interfaces to an area, and is flooded throughout a single area only.
- **Network LSA—Type-2 LSA.** Describes the list of routers connected to the network, and is flooded throughout a single area only.
- **Summary LSA—Type-3 LSA.** Describes the route to networks in another OSPF area of the same Autonomous System (AS). Propagated through backbone area to other areas.
- **ASBR summary LSA—Type-4 LSA.** Describes the route to an ASBR in an OSPF normal or backbone area of the same AS. Propagated through backbone area to other areas.
- **AS external LSA—Type-5 LSA.** Describes the route to a destination in another AS (external route.) Originated by ASBR in normal or backbone areas of an AS and propagates through backbone area to other normal areas.

- **NSSA LSA—Type-7 LSA.** Describes the route to a destination in another AS (external route.) Originated by ASBR in NSSA. ABR converts type-7 LSAs to type-5 LSAs for injection into the backbone area.

OSPFv2 router types

The following router types are supported

Internal routers

Internal OSPFv2 routers belong to only one area. Internal routers flood type-1 LSAs to all routers in the same area and maintain identical LSDBs.

Area border routers (ABRs)

Area border routers have membership in multiple areas. ABRs are used to connect the various areas in an AS to the backbone area for that AS. Multiple ABRs can be used to connect a given area to the backbone, and a given ABR can belong to multiple areas other than the backbone.

An ABR maintains a separate LSDB for each area to which it belongs. (All routers within the same area have identical LSDBs.) The ABR is responsible for flooding summary LSAs between its border areas. You can reduce summary LSA flooding by configuring area ranges. An area range enables you to assign an aggregate address to a range of IP addresses. This aggregate address is advertised instead of all the individual addresses it represents.

Autonomous system boundary router (ASBR)

Autonomous system boundary routers run multiple interior gateway protocols and serve as a gateway to other autonomous systems operating with interior gateway protocols. The ASBR imports and translates different protocol routes into OSPF through redistribution. ASBRs can be used in backbone areas, normal areas, and NSSAs, but not in stub areas.

Designated routers (DRs)

In an OSPF network having two or more routers, one router is elected to serve as the DR and another router to act as the Backup Designated Router (BDR). All other routers in the area forward their routing information to the DR and BDR, and the DR forwards this information to all routers in the network. This action minimizes the amount of repetitive information that is forwarded on the network by eliminating the need for each individual router in the area to forward its routing information to all other routers in the network. If the area includes multiple networks, each network elects its own DR and BDR.

In an OSPF network with no DR and no BDR, the neighboring router with the highest priority is elected the DR, and the router with the next highest priority is elected the BDR. If the DR goes off-line, the BDR automatically becomes the DR, and the router with the next highest priority then becomes the new BDR. If multiple routing switches on the same OSPF network are declaring themselves DRs, both priority and router ID are used to select the DR and BDRs.

Priority is configurable using the `ip ospf priority` command at the interface level. If two neighbors share the same priority, the router with the highest router ID is elected as the DR. The router with the next highest router ID is elected as the BDR.

OSPFv2 area types

OSPFv2 is built upon a hierarchy of network areas. All areas for a given OSPF domain reside in the same AS. An AS is defined as a number of contiguous networks that share the same interior gateway routing protocol.

An AS can be divided into multiple areas. Each area represents a collection of contiguous networks and hosts, and the topology of a given area is not known by the internal routers in any other area. Areas define the boundaries to which types 1 and 2 LSAs are broadcast, which limits the amount of LSA flooding that occurs within the AS and also helps to control the size of the LSDBs maintained in OSPF routers. An area is represented in OSPF by either an IP address or a number. Area types include: Backbone, Normal, Not-so-stubby (NSSA), and stub.

Backbone area

Every AS must have one (and only one) backbone area (identified as area 0 or 0.0.0.0.) The ABRs of all other areas in the same AS connect to the backbone area, either physically through an ABR or through a configured virtual link. The backbone is a transit area that carries the type-3 summary LSAs, type-5 AS external link LSAs and routed traffic between non-backbone areas, as well as the type-1 and type-2 LSAs and routed traffic internal to the area. ASBRs are allowed in backbone areas.

Normal area

Normal area connects to the backbone area through one or more ABRs (physically or through a virtual link) and supports type-3 summary LSAs and type-5 external link LSAs to and from the backbone area. ASBRs are allowed in normal areas.

Stub area

Stub area connects to the AS backbone through one or more ABRs. It does not allow an internal ASBR, and does not allow external (type 5) LSAs. A stub area supports these actions:

- Advertise the area summary routes to the backbone area.
- Advertise summary routes from other areas.
- Use the default summary (type-3) route to advertise both of the following:
  - Summary routes to other areas in the AS
  - External routes to other ASs

You can configure the stub area ABR to do the following:

- Suppress advertising from some or all area summarized internal routes into the backbone area.
- Suppress LSA traffic from other areas in the AS by replacing type-3 summary LSAs and the default external route from the backbone area with the default summary route (0.0.0.0/0.)

Virtual links are not allowed for stub areas.

Not-so-stubby (NSSA) area

NSSA area connects to the backbone area through one or more ABRs. NSSAs are intended for use where an ASBR exists in an area where you want to control the following:

- Advertising the ASBR external route paths to the backbone area
- Allowing LSAs from the backbone area to advertise in the NSSA:
  - Summary routes (type-3 LSAs) from other areas
  - External routes (type-5 LSAs) from other areas as a default external route (type-7 LSAs)

Virtual links are not allowed for NSSAs.

OSPFv2 configuration task list

Tasks at a glance

- Configuring OSPF on the routing switch
- Assigning the routing switch to an OSPF area
- Setting OSPF network for the area
- Creating an OSPF virtual link for an area
- Configuring external route redistribution and control
- Configuring area ranges on an ABR to reduce advertisements to the backbone
- Influencing route choice by changing the administrative distance
- Configuring graceful restart of OSPF routing
- Configuring OSPF interface settings
- Configuring OSPF interface authentication
- Configuring OSPF virtual link settings
- Configuring OSPF authentication on a virtual link
- Configuring all OSPF interfaces as passive
- Viewing OSPFv2 information
- Clearing OSPF statistics on a switch

Configuring OSPF on the routing switch

Create the OSPF instance and enter the OSPF router configuration context. From this, you can proceed with other OSPF configuration tasks.

Prerequisites

- You must be in the global configuration context, as indicated by the `switch(config)#` prompt to create the OSPF instance and enter the OSPF router configuration context.
- To configure a router ID, create OSPF network areas, or adjust other global OSPF configuration items, you must be in the router configuration context, as indicated by the `switch(config-router)#` prompt.

Procedure

1. Create the OSPF instance and enter the OSPF router configuration context using the following command. For command details, see router ospf.

   `router ospf <PROCESS-ID> [vrf <VRF-NAME>]`

   For example, the following command creates OSPF instance 1.

   ```
   switch(config)# router ospf 1
   switch(config-ospf-1)#
   ```

2. Configure a global router ID using the following command. For command details, see router-id .

   `router-id <ROUTER-ADDRESS>`

   For example, the following command sets the router ID to 1.1.1.1.

   ```
   switch(config-ospf-1)# router-id 1.1.1.1
   ```

3. Optionally, if the OSPF process was disabled (it is enabled by default), enable the OSPF process using the following command.
   `enable`
   For command details, see enable . (Refer to disable for disabling the OSPF process).

Assigning the routing switch to an OSPF area

Create a Normal, Stub, or Not So Stubby (NSSA) area.

Prerequisites

You must be in the router configuration context, as indicated by the `switch(config-router)#` prompt.

Procedure

Create an OSPF area for the routing switch using one of the following commands:

- Create a normal area using the following command: `area <area-id>`. For command details, see area (ospf).
- Create a stub area using the following command: `area <area-id> stub`. For command details, see area stub .
- Create a Not-So-Stubby-Area using the command: `area <area-id> nssa`. For command details, see area nssa .

For example, the following command creates a normal area with an area identifier of 10.1.1.1. Area identifier could alternatively be entered in decimal format such as 1.

```
switch(config-ospf-1)# area 10.1.1.1
```

Setting OSPF network for the area

After you define an OSPF area, you can assign one or more networks to it. The OSPF protocol will run on the interface with the configured IPv4 address. The interfaces which have an IP address configured in this network or in a subset of this network, will participate in the OSPF protocol.

Prerequisites

You must be in the interface configuration context, as indicated by the `switch(config-if)#` prompt.

Procedure

1. Set an OSPF network for the area using the following command. For command details, see ip ospf area.
   `ip ospf <PROCESS-ID> area <AREA-ID>`
   For example, use the following command to assign interface 1/1/1 to OSPF area 1. The area can alternatively be entered as an IPv4 address.

   ```
   switch(config)# interface 1/1/1
   switch(config-if)# ip ospf 1 area 1
   ```

2. Optionally, you can disable OSPF on the interface using the following command. For command details, see ip ospf shutdown.
   `ip ospf shutdown`

```
switch(config)# interface 1/1/1
switch(config-if)# ip ospf shutdown
```

Creating an OSPF virtual link for an area

Create an OSPF virtual link with remote ABR (if not created already) and enter the vlink context.

Prerequisites

You must be in the router configuration context, as indicated by the `switch(config-router)#` prompt.

Procedure

Create an OSPF virtual link using the following command. For command details, see area virtual-link .
`area <area-id> virtual-link <router-id>`

For example, the following command creates a virtual link in area 100.

```
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)#
```

Configuring external route redistribution and control

Configuring route redistribution for OSPFv2 establishes the routing switch as an ASBR for importing and translating different protocol routes into OSPFv2. When you configure redistribution for OSPFv2, you can specify that routes external to the OSPFv2 domain are imported as OSPFv2 routes.

1. Enable route redistribution using the `redistribute` command.
   `redistribute {bgp | connected | local loopback | static | rip | ospf <PROCESS-ID>} [route-map <ROUTE-MAP-NAME>]`

   For example, setting redistribution of connected routes as OSPFv2 routes:

```
switch(config-ospf-1)# redistribute connected
```

   If a route map is specified, then only the routes that pass the match clause specified in the route map are redistributed to OSPFv2. For example, setting redistribution of local loopback routes, that only match the routes specified by the route map:

```
switch(config-ospf-1)# redistribute local loopback route-map local_routes
```

2. Optionally, modify the default metric for redistribution using the `default-metric` command.
   `default-metric <METRIC-VALUE>`

   For example, setting the default metric for redistribution to 37.

```
switch(config-ospf-1)# default-metric 37
```

3. Optionally, set the cost of default-summary LSAs using the area command.
   `area <AREA-ID> default-metric <COST>`
   For example, setting the cost of default summary LSAs to 2.

```
switch (config-ospf-1)# area 1 default-metric 2
```

4. Optionally, use the `max-metric router-lsa` command to set the protocol to advertise a maximum metric so that other routers do not prefer this router as an intermediate hop in their shortest path first (SPF) calculations.
   `max-metric router-lsa [on-startup [<ADVERT-TIME>]]`
   For example, setting advertise max-metric router-lsa on startup.

```
switch(config-ospf-1)# max-metric router-lsa on-startup 3000
```

5. Optionally set the maximum number of ECMP routes that OSPFv2 can support using the `maximum-paths` command.
   `maximum-paths <MAX-VALUE>`
   For example, setting the maximum number of ECMP routes to 8.

```
switch(config-ospf-1)# maximum-paths 8
```

Configuring area ranges on an ABR to reduce advertisements to the backbone

You can configure area ranges to reduce inter-area advertisements by summarizing a range of IP addresses into a single route advertisement. This action prevents an ABR from advertising specific networks or subnets to the backbone area.

Prerequisites

You must be in the router configuration context, as indicated by the `switch(config-router)#` prompt.

Procedure

Summarize inter-area or NSSA paths using the following command. For command details, see area range .

`area <area-id> range <ip-prefix> type {inter-area | nssa} [no-advertise]`

For example, use the following command to summarize routes matching the area range 172.77.114.0/24 using inter-area as the type of address aggregation.

```
switch(config-ospf-1)# area 1 range 172.77.114.0/24 type inter-area
```

In another example, use the following command to specify DoNotAdvertise status for routes matching the area range 172.77.114.0/24. Use nssa as the type of address aggregation.

```
switch(config-ospf-1)# area 1 range 172.77.114.0/24 type nssa no-advertise
```

Influencing route choice by changing the administrative distance

The administrative distance value can be left in its default configuration setting (110) unless a change is needed to improve OSPF performance for a specific network configuration.

Prerequisites

You must be in the router configuration context, as indicated by the `switch(config-router)#` prompt.

Procedure

Reconfigure the administrative distance using the following command. For command details, see distance command.

`distance <distance>`

For example, use the following command to set administrative distance to 100.

```
switch(config-ospf-1)# distance 100
```

Configuring graceful restart of OSPF routing

OSPF routing can be gracefully restarted on the switch without losing packets that are in transit. OSPF neighbors are informed that the router is completing a graceful restart, which allows for maintenance on the switch without interrupting traffic in the network. There is no effect on the saved switch configuration.

Prerequisites

Graceful restart is only applicable to the 6300-VSF and 6400 switches.

You must be in the router configuration context, as indicated by the `switch(config-router)#` prompt.

Manually create a router-id configuration to support graceful restart and management module failover. A dynamic router-id is chosen from the list of configured IPs at the time of the router-id election, and in the event of a graceful restart and management module failover, the router-id should be preserved before and after these events. Graceful restart would fail if there is a change in the router-id, causing traffic loss.

Procedure

Enable graceful restart of OSPF routing using the following command. For command details, see graceful-restart .

`graceful-restart {restart-interval <seconds> | helper}`

For example, the following command specifies 50 seconds as the maximum interval another router will wait for this router to gracefully restart.

```
switch(config-ospf-1)# graceful-restart restart-interval 50
```

Configuring OSPF interface settings

You can optionally adjust the following OSPF interface settings.

Prerequisites

You must be in the interface configuration context, as indicated by the `switch(config-if)#` prompt.

Procedure

1. Set the interface cost using the following command. For command details, see ip ospf cost.
   `ip ospf cost <interface-cost>`
   For example, the following command sets the cost associated with interface 1/1/1 to 100.

   ```
   switch(config)# interface 1/1/1
   switch(config-if)# ip ospf cost 100
   ```

2. Set the time interval between OSPF hello packets for the OSPF interface using the following command. For command details, see ip ospf hello-interval.
   `ip ospf hello-interval <seconds>`
3. Set the interval after which a neighbor is declared dead if no hello packet is received on the OSPF interface using the following command. For command details, see ip ospf dead-interval.
   `ip ospf dead-interval <seconds>`
4. Set the time between retransmitting lost link state advertisements for the OSPF interface using the following command. For command details, see ip ospf retransmit-interval.
   `ip ospf retransmit-interval <seconds>`
5. Sets the transit delay in link state transmission for the OSPF interface using the following command. For command details, see ip ospf transit-delay.
   `ip ospf transit-delay <seconds>`
6. Set the OSPF network type for the interface. For command details, see ip ospf network.
   `ip ospf network {broadcast|point-to-point}`
7. Set the OSPF priority on the interface using the following command. For command details, see ip ospf priority.
   `ip ospf priority <number-value>`
8. Set the OSPF interface as OSPF passive interface. The interface participates in OSPF but does not send or receive packets on that interface. For command details, see ip ospf passive.
   `ip ospf passive`

Configuring OSPF interface authentication

Configure authentication on the interface. Only one method of authentication can be active on an interface at a time. If one method is already configured on an interface, configuring an alternative method on the same interface automatically overwrites the first method used.

Prerequisites

You must be in the interface configuration context, as indicated by the `switch(config-if)#` prompt.

Procedure

1. Set the authentication type that will be used for authentication with the neighbor router using the following command. For command details, see ip ospf authentication.

`ip ospf authentication {message-digest | simple-text | null | keychain}`

For example, the following command sets the authentication type to simple-text.

```
switch(config-if)# ip ospf authentication simple-text
```

2. For simple-text authentication, set the password using the following command. For command details, see ip ospf authentication-key.

```
ip ospf authentication-key <PASSWORD>
```

For example:

```
switch(config-if)# ip ospf authentication-key secure
```

3. For MD5 authentication, set the password using the following command. For command details, see ip ospf message-digest-key md5.

```
ip ospf message-digest-key {Key ID} md5 {ciphertext | plaintext}  <KEY>
```

For example:

```
switch(config-if)# ip ospf message-digest-key 1 md5 ciphertext 100
```

Configuring OSPF virtual link settings

The OSPF interface parameters for this process are automatically set to their default values for virtual links. You can optionally adjust the following OSPF virtual link settings.

Prerequisites

You must be in the router vlink configuration context, as indicated by the `switch(config-router-vlink)#` prompt.

1. Set the time interval between OSPF hello packets for the OSPF virtual links using the following command. For command details, see hello-interval.

```
hello-interval <SECONDS>
```

For example:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-ospf-1)# hello-interval 30
```

2. Set the interval which a neighbor is declared dead if no hello packet is received on the OSPF virtual links. Use the following command. For command details, see dead-interval.

```
dead-interval <SECONDS>
```

3. Set the time between retransmitting lost link state advertisements for the OSPF virtual links using the following command. For command details, see retransmit-interval.

```
retransmit-interval <SECONDS>
```

4. Sets the transit delay in Link state transmission for the OSPF virtual links using the following command. For command details, see transit-delay

```
transit-delay <SECONDS>
```

Configuring OSPF authentication on a virtual link

Set the OSPF virtual-link authentication type that will be used for authentication with the remote ABR when the command is used in the vlink context.

OSPF supports the same methods of authentication for virtual links as it does for interfaces in an area. In the default configuration, OSPF authentication is disabled. Only one method of authentication can be active on a virtual link at a time. If one method is configured on a virtual link, configuring the alternative method on the same link automatically replaces the first method. Both ends of a virtual link must use the same authentication method and related credentials.

Prerequisites

You must be in the router vlink configuration context, as indicated by the `switch(config-router-vlink)#` prompt.

Procedure

1. Set the OSPF virtual-link authentication type using the following command. For command details, see authentication.
   `authentication {message-digest | simple-text | null}`
   For example, the following command sets virtual-link authentication type to simple-text.

   ```
   switch(config-router-vlink)# authentication simple-text
   ```

2. For simple-text authentication, set the password using the following command. For command details, see authentication-key.
   `authentication-key <password>`
   `For example:`

   ```
   switch(config-router-vlink)# authentication-key plaintext secure
   ```

3. For MD5 authentication, set the password using the following command. For command details, see message-digest-key md5.
   `message-digest-key md5 <key>`
   `For example:`

   ```
   switch(config-router-vlink)# message-digest-key 1 md5 plaintext secure
   ```

Configuring all OSPF interfaces as passive

OSPF sends LSAs to all other routers in the same AS. To limit the flooding of LSAs throughout the AS, you can configure OSPF to be passive.

Prerequisites

You must be in the router configuration context, as indicated by the `switch(config-router)#` prompt.

Procedure

Configure all OSPF interfaces as passive using the following command. For command details, see passive-interface default.

```
passive-interface default
```

```
switch(config-ospf-1)# passive-interface default
```

Configuring SPF throttling timers

SPF calculation is throttled with default timer values (start-time 200ms, hold-time 1000ms, max-wait-time 5000ms). You can throttle SPF calculation by configuring non-default timers to improve performance of a specific network configuration.

Prerequisites

You must be in the router configuration context, as indicated by the `switch(config-router)#` prompt.

Procedure

Reconfigure SPF throttling timers using the following command. For command details, see timers throttle spf.

`timers throttle spf start-time <milliseconds> hold-time <milliseconds> max-wait-time <milliseconds>`

For example, use the following command to set start-time to 500ms, hold-time to 3000ms and max-wait-time to 9000ms:

```
switch(config-ospf-1)# timers throttle spf start-time 500 hold-time
3000 max-wait-time 9000
```

Viewing OSPFv2 information

Use these show commands from the Manager context, as indicated by the `switch#` prompt.

To view OSPFv2 information, use the following commands. For command details and examples, click the links.

- To view general OSPF, area, state and configuration information use: show ip ospf.
- To view OSPF routing table entries for Area Border Router (ABR) and Autonomous System Border Router (ASBR) use: show ip ospf border-routers.
- To view OSPF link state database summary for different OSPF LSAs (Link State Advertisement) use: show ip ospf lsdb.

Many different parameters are available with this command to display information for a particular LSA.

- To view information about OSPF-enabled interfaces use: show ip ospf interface.
- To view information about OSPF neighbors use: show ip ospf neighbors.
- To view OSPF routing table information use: show ip ospf routes .
- To view OSPF statistics use: show ip ospf statistics.
- To view OSPF statistics for the OSPF-enabled interfaces use show ip ospf statistics interface.
- To view the current state and parameters of the OSPF virtual-links use show ip ospf virtual-links.

Clearing OSPF statistics on a switch

Clear OSPF event statics using the following command. For command details, see clear ip ospf statistics .

```
clear ip ospf [<PROCESS-ID>] statistics [interface [<INTERFACE-NAME>]] [all-vrfs | vrf
<VRF-NAME>]
```

For example, the following command clears OSPF event statistics from interface 1/1/1.

```
switch(config-router)# clear ip ospf statistics interface 1/1/1
```

An example of the OSPFv2 information in the show running-config command

When OSPFv2 is configured, the output of the `show running-config` command includes OSPFv2 information.

For example:

```
switch# show running-config
Current configuration:
!
!Version AOS-CX 10.0X.XXXX
!
lldp enable
timezone set utc
vrf red
mgmt
led base-loc_fdc on
led base-loc on
led base-hlth_fdc fast_blink
led base-pwr_fdc on
!
!
!
!
!
!
aaa authentication login default local
aaa authorization commands default none
!
!
!
!
router ospf 1 vrf vrf_default
    area 0.0.0.0
    area 0.0.0.1
    area 0.0.0.1 virtual-link 55.4.4.4
        hello-interval 15
        dead-interval 60
router ospf 1 vrf red
    area 0.0.0.4
    area 0.0.0.5
    area 0.0.0.5 range 10.1.0.0/16 type inter-area
vlan 1
    no shutdown
interface lag 44
    no shutdown
    ip address 33.1.1.1/24
    ip ospf 1 area 0.0.0.0
    ip ospf hello-interval 22
interface 1/1/1
    no shutdown
```

```
    ip address 33.44.1.1/24
    ip ospf 1 area 0.0.0.0
    ip ospf passive
    ip ospf dead-interval 88
    ip ospf transit-delay 44
    ip ospf priority 88
    ip ospf cost 8
    ip ospf network point-to-point
    ip ospf authentication simple-text
    ip ospf authentication-key kkk
    ip ospf shutdown
interface 1/1/2
interface loopback 2
    ip address 55.55.55.55/32
    ip ospf 1 area 0.0.0.0
```

OSPFv2 commands

active-backbone

```
active-backbone stub-default-route
no active-backbone stub-default-route
```

Description

This command enables the router to send a default route to stub areas if there is an active loopback link in the backbone area. The configuration is not required if backbone area has neighbors or passive interfaces configured. By default active backbone detection is enabled.

Examples

```
switch(config)# router ospf 1
switch(config-ospf-1)# active-backbone stub-default-route
```

```
switch(config)# no active-backbone stub-default-route
```

Command History

Release	Modification
10.10.1000	Command Introduced

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325	`config-ospf-<PROCESS-ID>` `config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325H 8325P 8360 9300 9300S 10000		

area (ospf)

```
area <AREA-ID>
no area <AREA-ID>
```

Description

Creates a normal area, with **<AREA-ID>** set if not present. If the area is already present and it is not a normal area, then this command changes the area type to normal.

The **no** form of this command deletes the area with the **<AREA-ID>** specified. Area can be of any type (nssa, nssa no-summary, stub, stub no-summary, and default normal area).

Parameter	Description
<AREA-ID>	Specifies the area ID in one of the following formats. OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to 4294967295.

Examples

Creating a normal area:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 1
switch(config-ospf-1)# area 10.1.1.1
Switch(config-ospf-1)# show running-config current-context router ospf 1
       router-id 1.1.1.1
       area 0.0.0.0
       area 0.0.0.1
       area 0.0.0.2 stub
       area 0.0.0.3 nssa
```

Deleting an area:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no area 1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospf-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

area default-metric

```
area <AERA-ID> default-metric <COST>
no area <AREA-ID> default-metric
```

Description

Sets the cost of the default route announced to NSSA or stub areas.

The **no** form of this command resets the cost of the default route announced to NSSA or stub areas, to the default value of 1.

Parameter	Description
<AREA-ID>	Specifies area ID in one of the following formats. OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to 4294967295.
`default-metric <COST>`	Sets the cost of default-summary LSAs announced to NSSA or stub areas, to the specified value. Default cost: 1. Range: 0 to 16777215.

Examples

Setting cost for default LSA summary:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 1 default-metric 2
switch(config-ospf-1)# area 0.0.0.1 default-metric 2
```

Setting cost for default LSA summary to default:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no area 1 default-metric
switch(config-ospf-1)# no area 0.0.0.1 default-metric
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospf-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

area nssa

```
area <AREA-ID> nssa [no-summary]
no area <AREA-ID> nssa [no-summary]
```

Description

Creates the NSSA area (Not So Stubby Area) with **<AREA-ID>** if not present. If area is present and not NSSA area, this command changes the area type to NSSA area. If **no-summary** is used, area type will be NSSA No-Summary.

The **no** form of this command unsets the area type as NSSA. That is, the configured area will be changed to default normal area. The **no area <AREA-ID> nssa no-summary** command enables sending inter-area routes into NSSA, but will not unset the area as NSSA.

Parameter	Description
<AREA-ID>	Specifies the area ID in one of the following formats. OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to 4294967295.
`nssa [no-summary]`	Specifies Not So Stubby Area (NSSA) area type. If area is present and not NSSA area, parameter changes the area type to NSSA area. If **no-summary** is specified, area type will be NSSA No-Summary, which means do not inject inter-area routes into NSSA.

Examples

Creating an NSSA area:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 1 nssa
switch(config-ospf-1)# area 1 nssa no-summary
```

Unsetting the area as NSSA

```
switch(config)# router ospf 1
switch(config-ospf-1)# no area 1 nssa
switch(config-ospf-1)# no area 1 nssa no-summary
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospf-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

area range

```
area <AREA-ID> range <IP-PREFIX> type {inter-area | nssa} [no-advertise]
no area <AREA-ID> range <IP-PREFIX> type {inter-area | nssa} [no-advertise]
```

Description

Summarizes the routes with the matching address or masks. This command only works for border routers.

The **no** form of this command removes route summarization for the configured IPv4 prefix address on the ABR. When using the **no** form of the command with the **no-advertise** option, enables advertising this range to other areas.

Parameter	Description
<AREA-ID>	Specifies the area ID in one of the following formats. OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to

Parameter	Description
	4294967295.
`range <IP-PREFIX>`	Specifies summarizing routes matching the area range prefix/mask.
`type {inter-area | nssa}`	Specifies the type this address aggregation applies to as either inter-area range prefix or NSSA range prefix.
`no-advertise`	Specifies the address range status as **DoNotAdvertise** (do not advertise this range to other areas).

Examples

Summarizing inter-area or NSSA paths:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 1
switch(config-ospf-1)# area 2 nssa
switch(config-ospf-1)# area 1 range 192.77.114.0/24 type inter-area
switch(config-ospf-1)# area 2 range 192.77.114.0/24 type nssa
switch(config-ospf-1)# area 2 range 192.77.114.0/24 type nssa no-advertise
```

Removing summarization:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no area 1 range 192.77.114.0/24 type inter-area
switch(config-ospf-1)# no area 2 range 192.77.114.0/24 type nssa
switch(config-ospf-1)# no area 2 range 192.77.114.0/24 type nssa no-advertise
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospf-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

area stub

```
area <AREA-ID> stub [no-summary]
no area <AREA-ID> stub [no-summary]
```

Description

Creates the stub area with **<AREA-ID>** if not present. If the area is already present and it is not a normal stub area, then this command changes the stub area type to normal. If the **no-summary** parameter is used, area type will be stub No-Summary.

The **no** form of this command unsets the area as a stub type. That is, the configured area will be changed to a default normal area. The **no area <AREA-ID> stub no_summary** command enables sending inter-area routes into the stub area, but will not unset the area as stub.

ABR does not inject the default route in a Totally Stubby Area with loopback in Area 0.0.0.0. As a workaround, configure a passive interface or active neighbors in the backbone area.

Parameter	Description
<AREA-ID>	Specifies the area ID in one of the following formats. OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to 4294967295.
`stub [no-summary]`	Specifies the stub area type. If the area is already present and it is not a stub area, this parameter changes the area type to stub. If **no-summary** is specified, area type will be stub No-Summary (totally stubby area), which means do not inject summary link advertisements into stub areas.

Examples

Creating a STUB area:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 1 stub
switch(config-ospf-1)# area 1 stub no-summary
```

Unsetting the area type as stub:

```
switch(config)# router ospf 1
switch(config-ospf-1) # no area 1 stub
switch(config-ospf-1) # no area 1 sub no-summary
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospf-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

area virtual-link

```
area <AREA-ID> virtual-link <ROUTER-ID>
no area <AREA-ID> virtual-link <ROUTER-ID>
```

Description

Creates an OSPF virtual link with a remote ABR and enters the vlink context.

The **no** form of this command deletes an OSPF virtual link with the specified router ID of the remote ABR. If no **<ROUTER-ID>** is specified, the **no** form of the command sets the virtual link to the default settings.

Parameter	Description
`<AREA-ID>`	Specifies the area ID in one of the following formats. OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to 4294967295.
`virtual-link <ROUTER-ID>`	Configures a virtual link with the specified router ID of the remote ABR.

Examples

Configuring OSPF virtual links:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)#
```

Deleting OSPF virtual links:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no area 100 virtual-link 100.0.1.1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospf-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

authentication

```
authentication {hmac-sha-1 | hmac-sha-256 | hmac-sha-384 | hmac-sha-512 | message-digest
| simple-text | null | keychain}
no authentication
```

Description

Sets the OSPF virtual-link authentication type that will be used for authentication with the remote ABR.

Choose one of the authentication types from the following parameters.

The **no** form of this command unconfigures the virtual-link authentication type used and sets it to Null authentication.

Parameter	Description
`hmac-sha-1`	Sets the authentication type as SHA-1.
`hmac-sha-256`	Sets the authentication type as SHA-256.
`hmac-sha-384`	Sets the authentication type as SHA-384.
`hmac-sha-512`	Sets the authentication type as SHA-512.
`message-digest`	Sets the authentication type to message-digest.
`simple-text`	Sets the authentication type to simple-text.
`null`	Sets the authentication type to null.
`keychain`	Sets authentication type to use the key chain.

Examples

Setting OSPF virtual links authentication type:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# authentication simple-text
```

Deleting OSPF virtual links authentication type:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# no authentication
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-router-vlink`	Administrators or local user group members with execution rights for this command.

authentication-key

```
authentication-key [{ciphertext | plaintext} <PASSWORD>]
no authentication-key
```

Description

Sets the OSPF virtual-link authentication password that is used for simple-text authentication. If the password is given in ciphertext, it will be decrypted and applied to the protocol.

The **no** form of this command deletes the virtual-link authentication password that is used for simple-text authentication.

Parameter	Description
`{ciphertext | plaintext}`	Selects the password format.
<PASSWORD>	Specifies the password.

When the password is not provided on the command line, plaintext password prompting occurs upon pressing Enter. The entered password characters are masked with asterisks.

Examples

Setting the OSPF virtual link simple-text authentication password in plaintext format:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# authentication-key plaintext F82#450b
```

Setting the OSPF virtual link simple-text authentication with a prompted plaintext password:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# authentication-key
Enter the authentication key: ********
Re-Enter the authentication key: ********
```

Setting the OSPF virtual link simple-text authentication password in ciphertext format:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# authentication-key ciphertext AQaAz05...RmH+4pg=
```

Deleting the OSPF virtual link simple-text authentication password:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# no authentication-key
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360	config-router-vlink	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
9300 9300S 10000		

clear ip ospf neighbors

```
clear ip ospf [<PROCESS-ID>] neighbor [<NEIGHBOR>] [interface [<INTERFACE-NAME>]]
[all-vrfs | vrf <VRF-NAME>]
```

Description

Resets the neighbor and clears the OSPF neighbor information.

Parameter	Description
<PROCESS-ID>	Specifies the OSPFv2 process ID to clear the statistics for the particular OSPFv2 process. Range: 1 to 65535.
<NEIGHBOR>	Specifies the router ID of a neighbor.
<INTERFACE-NAME>	Specifies the OSPFv2 statistics to clear for the specified interface.
`all-vrfs`	Select to clear the OSPFv2 statistics for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF.

Example

Clearing the OSPFv2 neighbor information:

```
switch# clear ip ospf 1 neighbor
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ip ospf 1 neighbor 1.1.1.2
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ip ospf 1 neighbor interface 1/1/1
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ip ospf 1 neighbor 1.1.1.5 vrf red
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ip ospf neighbor
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ip ospf neighbor 1.1.1.4
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ip ospf neighbor interface 1/1/1
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ip ospf neighbor 1.1.1.5 vrf red
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

clear ip ospf statistics

`clear ip ospf [<PROCESS-ID>] statistics [interface [<INTERFACE-NAME>]] [all-vrfs | vrf <VRF-NAME>]`

Description

Clear the OSPF event statistics.

Parameter	Description
<PROCESS-ID>	OSPF process ID. Clear the statistics for the particular OSPF process. Range: 1 to 65535.
<INTERFACE-NAME>	Clear the OSPF statistics for the specified interface.
`all-vrfs`	Optionally select to clear the OSPF statistics for all VRFs.
`vrf` *<VRF-NAME>*	Optionally select to clear the OSPF statistics for a particular VRF. If the VRF is not specified, information for the default VRF is cleared.

Examples

Clearing the OSPF event statistics:

```
switch# clear ip ospf statistics
switch# clear ip ospf statistics interface 1/1/1
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

dead-interval

```
dead-interval <INTERVAL>
no dead-interval
```

Description

Sets the interval after which a neighbor is declared dead if no hello packet comes in for virtual links.

The **no** form of this command sets the dead interval to default for virtual links. The default value is 40 seconds (generally four times the hello packet interval).

For proper operation, set the dead interval must be longer than the hello interval.

Parameter	Description
<INTERVAL>	Specifies the time interval for the dead interval, in seconds. Range: 1 to 65535. Default: 40.

Examples

Setting the OSPv2F virtual links dead interval:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# dead-interval 30
```

Setting the OSPFv2 virtual links dead interval to default:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# no dead-interval
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-router-vlink`	Administrators or local user group members with execution rights for this command.

default-information originate

```
default-information originate [metric <METRIC-VALUE>]
no default-information originate [metric <METRIC-VALUE>]
```

Description

Configures OSPF to advertise the default route (0.0.0.0/0) to its neighbors if it is present in the routing table. Optionally, the metric value can be set for default route ::/0. The default value is 1.

The **no** form of this command disables advertisement of the default route.

Parameter	Description
metric <METRIC-VALUE>	Specifies the OSPF metric value for the default route. Optional. Default: 1.

Examples

Setting advertisement of the default route:

```
switch(config)# router ospf 1
switch(config-ospf-1)# default-information originate
```

Disabling advertisement of the default route:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no default-information originate
```

Setting advertisement of the default route and specifying an optional metric value of 20:

```
switch(config)# router ospf 1
switch(config-ospfv3-1)# default-information originate
switch(config-ospfv3-1)# default-information originate metric 20
```

Disabling advertisement of the default route and setting metric to the default value:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no default-information originate metric
```

Command History

Release	Modification
10.09	Added parameter: **metric** <**METRIC-VALUE**>
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospf-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

default-information originate always

```
default-information originate always [metric <METRIC-VALUE>]
no default-information originate always [metric <METRIC-VALUE>]
```

Description

Configures OSPF to advertise the default route (0.0.0.0/0) to its neighbors, regardless if it is present in the routing table or not. Optionally, metric can be set for default route 0.0.0.0/0. The default value is 1.

The **no** form of this command disables advertisement of the default route.

Parameter	Description
metric <METRIC-VALUE>	Specifies the OSPF metric value for the default route. Default: 1.

Examples

Setting advertisement of the default route:

```
switch(config)# router ospf 1
switch(config-ospf-1)# default-information originate always
```

Disabling advertisement of the default route:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no default-information originate always
```

Setting advertisement of the default route with metric set to 20:

```
switch(config)# router ospf 1
switch(config-ospf-1)# default-information originate always metric 20
```

Disabling advertisement of the default route and setting the metric to the default value:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no default-information originate always metric
```

Command History

Release	Modification
10.09	Added parameter: metric <**METRIC-VALUE**>
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospf-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

default-metric

```
default-metric <METRIC-VALUE>
no default-metric
```

Description

Sets the default metric for redistributed routes in the OSPF.

The **no** form of this command sets the default metric to be used for redistributed routes into OSPF to the default of 25.

Parameter	Description
<METRIC-VALUE>	Specifies the default metric value to use for redistributed routes. Default: 25. Range: 0-1677214.

Examples

Setting default metric for redistributed routes:

```
switch(config)# router ospf 1
switch(config-ospf-1)# default-metric 37
```

Setting default metric for redistributed routes to the default:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no default-metric
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospf-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

disable

```
disable
```

Description

Disables the OSPF process.

This command does not remove the OSPF configurations.

Examples

Disabling OSPF process:

```
switch(config)# router ospf 1
switch(config-ospf-1)# disable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospf-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

distance

```
distance [<DISTANCE-VAL> | intra-area [<DISTANCE-VAL>] | inter-area [<DISTANCE-VAL>] |
external [<DISTANCE-VAL>]]
no distance [intra-area | inter-area | external]
```

Description

Defines an Administrative Distance (AD) for OSPF. Administrative Distance is used as a criteria to select the best route when multiple routes are present from different routing protocols.

The **no** form of this command sets the OSPF administrative distance to the default of 110. Optionally, administrative distance can be set to default for the specific OSPF route type: intra-area, inter-area, or external type-5 and type-7 routes.

Parameter	Description
<DISTANCE-VAL>	Specifies the OSPF administrative distance. Range: 1 to 255. Default: 110.
intra-area	Specifies the OSPF distance for intra-area routes.
inter-area	Specifies the OSPF distance for inter-area routes.
external	Specifies the OSPF distance for external type 5 and type 7 routes.

Usage

Within a given OSPF process, intra-area routes are always given precedence even when distances are configured for inter-area or external type routes.

Examples

Setting OSPF administrative distance:

```
switch(config)# router ospf 1
switch(config-ospf-1)# distance 100
switch(config-ospf-1)# distance intra-area 24 external 55 inter-area 66
switch(config-ospf-1)# distance intra-area 24 external 55
switch(config-ospf-1)# distance external 55
switch(config-ospf-1)#exit

switch(config)# router ospf 2
switch(config-ospf-2)# distance 200
switch(config-ospf-2)# distance external 60
switch(config-ospf-2)# distance intra-area 24 inter-area 66
```

Setting OSPF administrative distance to the default:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no distance
switch(config-ospf-1)# no distance external
switch(config-ospf-1)# no distance inter-area
switch(config-ospf-1)# no distance intra-area
switch(config-ospf-1)# no distance 100
switch(config-ospf-1)# no distance 220

switch(config)# router ospf 2 vrf blue
switch(config-ospf-2)# no distance 200
switch(config-ospf-2)# no distance external 60
switch(config-ospf-2)# no distance intra-area 24 inter-area 66
```

Command History

Release	Modification
10.14	Added capability to have individual admin distance for multiple OSPF processes in a VRF.
10.09	Added parameters: **intra-area**, **inter-area**, **external**
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100	`config-ospf-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325 8325H 8325P 8360 9300 9300S 10000		

distribute-list prefix

```
distribute-list prefix <prefix-list-name> {in | out}
no distribute-list prefix <prefix-list-name> {in | out}
```

Description

This command uses an existing prefix list to filter routes that are being installed in the routing table or redistributed to another routing protocol.

The **distribute-list prefix** command filters routes in the inbound or the outbound direction. When this command is issued with the **in** parameter, it filters routes from being installed in the routing table, it does not filter LSAs. When this command is issued with the **out** parameter, it filters only the desired redistributed routes from other protocols.

This command requires that your prefix list is already defined using the ip prefix commands. Route-maps are not supported with the distribute-list feature.

Parameter	Description
`prefix <prefix-list-name>`	Specify the name of an existing prefix.
`{in | out}`	Select one of the following parameters to set the filter direction: ▪ **in**: Filter incoming routes into the routing table ▪ **out**: Filter outgoing routing updates

Examples

The following commands enable the filtering of OSPFv2 routes in an IPv4 network, so routes are no longer installed in the routing table or redistributed from another routing protocol.

```
switch(config)# router ospfv2 1
switch(config-ospfv2-1)# distribute-list prefix listA in
switch(config-ospfv2-1)# distribute-list prefix listB out
```

The following command disables the filtering of OSPFv2 routes in an IPv4 network, so routes can be installed in the routing table or redistributed from another routing protocol.

```
switch# configure terminal
switch(config)# router ospfv2 1
switch(config-ospfv2-1)# no distribute-list prefix listA in
switch(config-ospfv2-1)# no distribute-list prefix listB out
```

Command History

Release	Modification
10.13	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv2-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

enable

`enable`

Description

Enables the OSPF process, if disabled. By default the OSPF process is enabled.

Examples

Enabling OSPF process:

```
switch(config)# router ospf 1
switch(config-ospf-1)# enable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H	`config-ospf-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325P 8360 9300 9300S 10000		

graceful-restart

```
graceful-restart
   restart-interval <INTERVAL>
   helper [strict-lsa-check]
   ignore-lost-interface
   no...
```

Description

Configures graceful restart parameters for OSPF.

The **no** form of this command sets the restart interval to the default interval of 120 seconds or disables the helper mode, depending on the parameter specified.

Parameter	Description
`restart-interval <INTERVAL>`	Specifies the time another router waits for this router to gracefully restart and selects the maximum time to wait in seconds. Range: 5 to 1800. Default: 120.
`helper`	Specifies that the router will participate in the graceful restart of a neighbor router.
`strict-lsa-check`	(Optional). Use with the **helper** parameter to enable strict Link state Advertisement (LSA) checking when acting as a restart helper for a restarting peer. **NOTE:** OSPF neighbors must disable strict LSA checking. If the local node has fewer OSPF interfaces after restarting, then the neighbors that were adjacent on those interfaces will clear up their adjacencies to the restarting node and will send out link state updates to advertise the dropped adjacency. If strict LSA checking is enabled, the restarting router's neighbors will exit helper mode when they receive the updated LSAs and the graceful restart will still fail.
`ignore-lost-interface`	Enable the restarting router to ignore lost OSPF interfaces during a graceful restart process. This setting should be enabled on a high availability system to ensure a graceful restart completes successfully, even if OSPF-enabled links fail due to High Availability events like a switchover or failover. **NOTE:** Enabling this setting means that the hitless restart procedures do not strictly follow those defined in *RFC 3623, Graceful OSPF Restart*.
`no`	Negate any parameter or return the setting to its default.

Examples

Enabling OSPF graceful restart:

```
switch(config)# router ospf 1
switch(config-ospf-1)# graceful-restart restart-interval 40
switch(config-ospf-1)# graceful-restart helper strict-lsa-check
```

Enabling the switch to ignore lost OSPF interfaces during a graceful restart process:

```
switch(config)# router ospf 1
switch (config-ospf-1)# graceful-restart ignore-lost-interface
```

Setting the restart interval to default, and disabling helper mode:

```
switch(config)# router ospf 1
switch(config-ospfv3-1)# no graceful-restart restart-interval
switch(config-ospfv3-1)# no graceful-restart helper
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospf-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

hello-interval

```
hello-interval <INTERVAL>
no hello-interval
```

Description

Sets the time interval between OSPF hello packets for virtual links.

The **no** form of this command sets the hello interval to the default value of 10 seconds for virtual links.

For proper operation, the hello interval must be shorter than the dead interval.

Parameter	Description
<INTERVAL>	Specifies the time interval for the hello interval, in seconds. Range: 1 to 65535. Default: 10.

Examples

Setting the OSPF virtual links hello interval:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# hello-interval 30
```

Setting the OSPF virtual links hello interval to default:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# no hello-interval
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-router-vlink`	Administrators or local user group members with execution rights for this command.

ip ospf area

```
ip ospf <PROCESS-ID> area <AREA-ID>
no ip ospf <PROCESS-ID> area <AREA-ID>
```

Description

Runs the OSPF protocol on the interface with the configured IPv4 address for the area specified. The interfaces which have an IP address configured in this network or in a subset of this network, will participate in the OSPF protocol.

To move an interface to a new area, unmap the existing area and then associate a new area with the interface.

The **no** form of this command disables OSPF on the interface and removes the interface from the area. Interfaces which have an IP address configured on the network or in a subset of the network, stop participating in the OSPF protocol.

Parameter	Description
<PROCESS-ID>	Specifies the OSPF process Id. Range: 1 to 65535.
<AREA-ID>	Specifies the OSPF area ID in one of the following formats. Area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. Area identifier in decimal format. Range: 0 to 4294967295.

Examples

Setting OSPF network for the area:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf 1 area 1
switch(config-if-vlan)# ip ospf 1 area 0.0.0.1
```

Disabling OSPF network for the area:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ip ospf 1 area 1
switch(config-if-vlan)# no ip ospf 1 area 0.0.0.1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

ip ospf authentication

```
ip ospf authentication {message-digest | simple-text | null | keychain | hmac-sha-1 |
hmac-sha-256 | hmac-sha-384 | hmac-sha-512}
no ip ospf authentication
```

Description

Sets the authentication type that will be used for authentication with the neighbor router.

The **no** form of this command deletes the authentication type used for a particular authentication with the neighbor router and sets to null authentication.

Parameter	Description
`message-digest`	Sets authentication type as message-digest.
`simple-text`	Sets authentication type as simple-text.
`null`	Sets authentication type as null.
`keychain`	Sets the authentication type to use the key chain.
`hmac-sha-1`	Sets the authentication type to SHA-1.
`hmac-sha-256`	Sets the authentication type to SHA-256.
`hmac-sha-384`	Sets the authentication type to SHA-384.
`hmac-sha-512`	Sets the authentication type to SHA-512.

Examples

Setting OSPF authentication type on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf authentication simple-text
```

Deleting OSPF authentication type on the interface and sets it to null:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ip ospf authentication
```

Setting OSPF authentication type to SHA-384 on the interface:

```
switch(config)# interface vlan 5
switch(config-if-vlan)# ip ospf authentication hmac-sha-384
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority

6300, except for S3L75A, S3L76A, S3L77A
6400
8100
8325
8325H
8325P
8360
9300
9300S
10000
	`config-if`
	`config-if-vlan`
	`config-lag-if`
		Administrators or local user group members with execution rights for this command.

ip ospf authentication-key

```
ip ospf authentication-key [{ciphertext | plaintext} <PASSWORD>]
no ip ospf authentication-key
```

Description

Sets the authentication password used for simple-text authentication. If the password is given in ciphertext it will be decrypted and applied to the protocol.

The **no** form of this command deletes the authentication password used for simple-text authentication.

Parameter	Description
`{ciphertext | plaintext}`	Selects the password format.
<PASSWORD>	Specifies the password.

When the password is not provided on the command line, plaintext password prompting occurs upon pressing Enter. The entered password characters are masked with asterisks.

Examples

Setting the OSPF simple-text authentication password in plaintext format:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf authentication-key plaintext F82#450b
```

Setting the OSPF simple-text authentication password with a prompted plaintext password:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf authentication-key
Enter the authentication key: ********
Re-Enter the authentication key: ********
```

Setting the OSPF simple-text authentication password in ciphertext format:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf authentication-key ciphertext AQBaZ...ecopg=
```

Deleting the OSPF simple-text authentication password:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ip ospf authentication-key
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

ip ospf cost

```
ip ospf cost <INTERFACE-COST>
no ip ospf cost
```

Description

Sets the cost (metric) associated with a particular interface. The interface cost is used as a parameter to calculate the best routes.

The **no** form of this command sets the cost (metric) associated with a particular interface to the default cost 1.

Parameter	Description
<INTERFACE-COST>	Specifies the interface cost value. Range: 1 to 65535. Default: 1.

Examples

Setting OSPF interface cost

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf cost 100
```

Setting the OSPF interface cost to default

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ip ospf cost
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

ip ospf dead-interval

```
ip ospf dead-interval <INTERVAL>
no ip ospf dead-interval
```

Description

Sets the interval after which a neighbor is declared dead if no hello packet is received on the OSPF interface.

The **no** form of this command sets the interval after which a neighbor is declared dead, to the default for the OSPF interface. The default value is 40 seconds (generally 4 times the hello packet interval).

Parameter	Description
<INTERVAL>	Specifies the time interval for the dead interval, in seconds. Range: 1 to 65535. Default: 40.

Examples

Setting OSPF dead interval on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf dead-interval 30
```

Setting OSPF dead interval to default on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ip ospf dead-interval
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

ip ospf hello-interval

```
ip ospf hello-interval <INTERVAL>
no ip ospf hello-interval
```

Description

Sets the time interval between OSPF hello packets for the OSPF interface.

The **no** form of this command sets the time interval OSPF hello packets to the default of 10 seconds for the OSPF interface.

Parameter	Description
<INTERVAL>	Specifies the time interval for the hello interval, in seconds. Range: 1 to 65535. Default: 10.

Examples

Setting OSPF hello interval on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf hello-interval 30
```

Setting OSPF hello interval to the default on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ip ospf hello-interval
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

ip ospf keychain

```
ip ospf keychain <KEYCHAIN-NAME>
no ip ospf keychain
```

Description

Sets the key chain for md5 authentication. A key chain configures rotating keys for packet authenticating, reducing the risk of keys being compromised.

The **no** form of this command deletes the key chain used for md5 authentication.

Parameter	Description
<KEYCHAIN-NAME>	Name of key chain to be used for md5 authentication.

Examples

Setting OSPFv2 key chain authentication:

```
switch(config)# interface 1/1/1
switch(config-if)# ip ospf keychain ospf_keys
```

Deleting OSPFv2 key chain authentication:

```
switch(config)# interface 1/1/1
switch(config-if)# no ip ospf keychain
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ip ospf message-digest-key md5

```
ip ospf message-digest-key <KEY-ID> md5 [{ciphertext | plaintext} <KEY>]
no ip ospf message-digest-key <KEY-ID>
```

Description

Sets the md5 message digest authentication key. If the md5 key is given in ciphertext, it will be decrypted and applied to the protocol.

The **no** form of this command deletes the md5 authentication key.

Parameter	Description
<KEY-ID>	Specifies the md5 key ID. Range: 1 to 255.
{ciphertext | plaintext}	Selects the md5 key format.
<KEY>	Specifies the md5 authentication key.

When the authentication key is not provided on the command line, plaintext key prompting occurs upon pressing Enter. The entered key characters are masked with asterisks.

Examples

Setting the md5 key in plaintext format:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf message-digest-key 1 md5 plaintext F82#450b
```

Setting the md5 key with a prompted plaintext key:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf message-digest-key 1 md5
Enter the MD5 authentication key: ********
Re-Enter the MD5 authentication key: ********
```

Setting the md5 key in ciphertext format:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf message-digest-key 1 md5 ciphertext AQt6e...7qEa4=
```

Deleting the md5 key:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ip ospf message-digest-key 1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

ip ospf network

```
ip ospf network {broadcast | point-to-point}
no ip ospf network
```

Description

Configures the OSPF network type for the interface. Choose one of the following parameters as the interface network type.

The **no** form of this command sets the network type for the interface to the system default which is broadcast network.

For an IP unnumbered interface, the network type should be configured as a **point-to-point**.

Parameter	Description
`broadcast`	Specifies the OSPF network type as a broadcast multi-access network.
`point-to-point`	Specifies the OSPF network type as a point-to-point network.

Examples

Setting OSPF network type for the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf network broadcast
switch(config-if-vlan)# ip ospf network point-to-point
```

Disabling OSPF network type for the interface to system default of broadcast network:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ip ospf network
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325P 8360 9300 9300S 10000		

ip ospf passive

```
ip ospf passive
no ip ospf passive
```

Description

Configures the interface as an OSPF passive interface. With this setting, the interface participates in OSPF but does not send or receive packets on that interface.

The **no** form of this command resets the interface as active. With this setting, the interface starts sending and receiving OSPF packets.

Examples

Setting the interface as OSPF passive interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf passive
```

Setting the interface as OSPF active interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ip ospf passive
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if config-if-vlan	Administrators or local user group members with execution rights for this command.

ip ospf priority

```
ip ospf priority <PRIORITY-VALUE>
no ip ospf priority
```

Description

Sets the OSPF priority for the interface. The larger the numeric value of the priority, the higher the chances for it to become the designated router. Setting a priority of zero makes the router ineligible to become a designated router or back up designated router.

The **no** form of this command sets the OSPF priority for the interface to the default of 1.

Parameter	Description
<PRIORITY-VALUE>	Specifies the OSPF priority value. Range: 0 to 255. Default: 1.

Examples

Setting OSPF priority for the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf priority 50
```

Disabling OSPF priority for the interface to default:

```
switch(config)# interface vlan1
switch(config-if-vlan)# no ip ospf priority
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if config-if-vlan	Administrators or local user group members with execution rights for this command.

ip ospf retransmit-interval

```
ip ospf retransmit-interval <INTERVAL>
```

```
no ip ospf retransmit-interval
```

Description

Sets the time between retransmitting lost link state advertisements for the OSPF interface.

The **no** form of this command sets the time between retransmitting lost link state advertisements to the default of 5 seconds for the OSPF interface.

Parameter	Description
<INTERVAL>	Specifies the retransmit interval, in seconds. Range: 1 to 3600. Default: 5.

Examples

Setting OSPF retransmit interval on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf retransmit-interval 30
```

Setting OSPF retransmit interval to the default on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ip ospf retransmit-interval
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

ip ospf sha-key sha

```
ip ospf sha-key <KEY-ID> sha [{ciphertext | plaintext} <KEY>]
no ip ospf sha-key <KEY-ID>
```

Description

Sets the SHA (secure hash authentication) key for the selected interface. If the SHA key is given in ciphertext, it will be decrypted and applied to the protocol. This command accepts a key of up to 64 characters irrespective of the SHA version configured on the interface. OSPF will internally pad zeros to the key to obtain a 64-byte key. For all types of SHA, key length is adjusted to 64 bytes.

The **no** form of this command deletes the SHA authentication key.

Parameter	Description
<KEY-ID>	Specifies the SHA key ID. Range: 1 to 255.
`{ciphertext | plaintext}`	Selects the SHA key format.
<KEY>	Specifies the SHA authentication key.

When the authentication key is not provided on the command line, plaintext key prompting occurs upon pressing Enter. The entered key characters are masked with asterisks.

Examples

Setting the SHA authentication key in plaintext format:

```
switch(config)# interface 1/1/1
switch(config-if)# ip ospf sha-key 1 sha plaintext F82#450b
```

Setting the SHA authentication key in prompted plaintext format:

```
switch(config)# interface 1/1/1
switch(config-if)# ip ospf sha-key 1 sha
Enter the SHA authentication key: ********
Re-Enter the SHA authentication key: ********
```

Setting the SHA authentication key in ciphertext format:

```
switch(config)# interface 1/1/1
switch(config-if)# ip ospf sha-key 1 sha ciphertext AQapu...C2K47A=
```

Deleting the SHA authentication key:

```
switch(config)# interface 1/1/1
switch(config-if)# no ip ospf sha-key 1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority

6300, except for S3L75A, S3L76A, S3L77A
6400
8100
8325
8325H
8325P
8360
9300
9300S
10000
	`config-if`
	`config-if-vlan`
	`config-lag-if`
		Administrators or local user group members with execution rights for this command.

ip ospf shutdown

```
ip ospf shutdown
no ip ospf shutdown
```

Description

Disables OSPF on the interface. The interface state changes to Down. It does not remove the interface from the OSPF area. To remove the interface, use the command **no ip ospf area**.

The **no** form of this command re-enables OSPF on the interface

Examples

Disabling OSPF on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf shutdown
```

Re-enabling OSPF on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ip ospf shutdown
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300 9300S 10000		

ip ospf transit-delay

```
ip ospf transit-delay <DELAY>
no ip ospf transit-delay
```

Description

Sets the time delay in link state transmission for the OSPF interface.

The **no** form of this command sets the delay in link state transmission to the default of 1 second for the OSPF interface.

Parameter	Description
<DELAY>	Specifies the transit delay in seconds. Range: 1 to 3600. Default: 1.

Examples

Setting OSPF transit delay on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ip ospf transit-delay 30
```

Setting OSPF transit delay to the default on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ip ospf transit-delay
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300 9300S 10000		

keychain

```
keychain <KEYCHAIN-NAME>
no keychain
```

Description

Sets the key chain for md5 authentication. A key chain configures rotating keys for packet authenticating, reducing the risk of keys being compromised.

The **no** form of this command deletes the key chain used for md5 authentication.

Parameter	Description
<KEYCHAIN-NAME>	Name of key chain to be used for md5 authentication.

Examples

Setting OSPF virtual link key chain authentication:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# keychain ospf_keys
```

Deleting OSPF virtual link key chain authentication:

```
switch# configure terminal
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link  100.0.1.1
switch(config-router-vlink)# no keychain
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except	config-router-vlink	Administrators or local user group members with execution

Platforms	Command context	Authority
for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000		rights for this command.

max-metric router-lsa

```
max-metric router-lsa [on-startup [<ADVERT-TIME>]]
no max-metric router-lsa [on-startup]
```

Description

Sets the protocol to advertise a maximum metric so that other routers do not prefer this router as an intermediate hop in their shortest path first (SPF) calculations. If the on-startup parameter is used, the router is configured to advertise a maximum metric at startup for the time mentioned in seconds or for a default value of 600 seconds.

The **no** form of this command advertises the normal cost metrics instead of advertising the maximized cost metric. This setting causes the router to be considered in traffic forwarding.

Parameter	Description
`on-startup <ADVERT-TIME>`	Specifies the time in seconds to advertise self as stub-router on startup. If no time is specified, the default time of 600 seconds is used. Range: 5 to 86400. Default: 600.

Examples

Setting to maximize the cost metrics for Router LSA:

```
switch(config)# router ospf 1
switch(config-ospf-1)# max-metric router-lsa
switch(config-ospf-1)# max-metric router-lsa on-startup
switch(config-ospf-1)# max-metric router-lsa on-startup 3000
```

Setting to advertise the normal cost metrics instead of advertising the maximized cost metric:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no max-metric router-lsa
switch(config-ospf-1)# no max-metric router-lsa on-startup
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospf-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

maximum-paths

```
maximum-paths <MAX-VALUE>
no maximum-paths
```

Description

Sets the maximum number of ECMP routes that OSPF can support.

The **no** form of this command sets the maximum number of ECMP routes that OSPF can support to the default value of 4.

Parameter	Description
<MAX-VALUE>	Specifies the maximum number of ECMP routes. Range: 1 to 32. Default: 4.

Examples

Setting maximum number of ECMP routes:

```
switch(config)# router ospf 1
switch(config-ospf-1)# maximum-paths 32
```

Setting maximum number of ECMP routes to the default of 4:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no maximum-paths
```

Command History

Release	Modification
10.10	Increased upper limit of range of **<MAX-VALUE>** parameter to 32.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospf-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

message-digest-key md5

```
message-digest-key <KEY-ID> md5 [{ciphertext | plaintext} <KEY>]
no message-digest-key <KEY-ID>
```

Description

Sets the virtual link md5 message digest authentication key. If the md5 key is given in ciphertext, it will be decrypted and applied to the protocol.

The **no** form of this command deletes the virtual link md5 authentication key.

Parameter	Description
<KEY-ID>	Specifies the virtual link md5 key ID. Range: 1 to 255.
`{ciphertext | plaintext}`	Selects the virtual link md5 key format.
<KEY>	Specifies the virtual link md5 authentication key.

When the authentication key is not provided on the command line, plaintext key prompting occurs upon pressing Enter. The entered key characters are masked with asterisks.

Examples

Setting virtual link md5 authentication key in plaintext format:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# message-digest-key 1 md5 plaintext F82#450b
```

Setting the virtual link md5 authentication key in prompted plaintext format:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# message-digest-key 1 md5
Enter the MD5 authentication key: ********
Re-Enter the MD5 authentication key: ********
```

Setting the virtual link md5 authentication key in ciphertext format:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# message-digest-key 1 md5 ciphertext AQapu...C2K47A=
```

Deleting the virtual link md5 authentication password:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# no message-digest-key 1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-router-vlink`	Administrators or local user group members with execution rights for this command.

passive-interface default

```
passive-interface default
no passive-interface
```

Description

Configures all OSPF interfaces as passive.

The **no** form of this command sets all OSPF interfaces as active.

Examples

Setting OSPF-enabled interfaces as passive:

```
switch(config)# router ospf 1
switch(config-ospf-1)# passive-interface default
```

Setting OSPF-enabled interfaces as active:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no passive-interface
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospf-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

redistribute

```
redistribute {bgp | connected | host-routes | local loopback | static | rip | ospf
<PROCESS-ID>[route-map <ROUTE-MAP-NAME>]
no redistribute {bgp | connected | host-routes | local loopback | static | rip | ospf
<PROCESS-ID>}[route-map <ROUTE-MAP-NAME>]
```

Description

Redistributes routes originating from other protocols, or from another OSPFv2 process, to the current OSPFv2 process.

If a route map is specified, then only the routes that pass the match clause specified in the route map are redistributed to OSPFv2. Configuration is not allowed if the referenced route map has not yet been configured.

If you try to redistribute routes from an OSPFv2 process which is not created, you are prompted to allow the OSPFv2 process to be auto-created before proceeding with redistribution. If you confirm at the prompt, the OSPFv2 process is created with defaults and redistribution configuration applied. If you deny at the prompt, redistribution configuration is skipped.

If command **route-redistribute active-routes-only** has been issued, only the routes from other protocols which are selected for forwarding are considered for redistribution into OSPFv2.

The **no** form of this command disables redistribution of routes to the current OSPFv2 process.

Parameter	Description
`bgp`	Specifies redistributing BGP (Border Gateway Protocol) routes.
`connected`	Specifies redistributing connected (directly attached subnet or host).
`host-routes`	Redistributes ARP entries. ARP entries learnt through EVPN or over VXLAN tunnels are not redistributed. The **host-routes** parameter is applicable to HPE Aruba Networking 8325, 10000, 8360, and 9300/9300S Switch series only.
`local loopback`	Specifies redistributing local routes of the loopback interface.
`static`	Specifies redistributing static routes.
`rip`	Specifies redistributing RIP routes.
`ospf <PROCESS-ID>`	Specifies redistributing routes from the specified OSPFv2 process ID. Range: 1 to 65535.
`route-map <ROUTE-MAP-NAME>`	Specifies redistribution filtering by route map. To create a route map, use command **route-map**.

Examples

Redistributing routes to OSPFv2:

```
switch(config)# router ospf 1
switch(config-ospf-1)# redistribute bgp
switch(config-ospf-1)# redistribute bgp route-map BGP_routes
```

```
switch(config-ospf-1)# redistribute host-routes
switch(config-ospf-1)# redistribute connected
switch(config-ospf-1)# redistribute connected route-map connected_routes
switch(config-ospf-1)# redistribute local loopback
switch(config-ospf-1)# redistribute local loopback route-map local_routes
switch(config-ospf-1)# redistribute static
switch(config-ospf-1)# redistribute static route-map static_networks
switch(config-ospf-1)# redistribute rip
switch(config-ospf-1)# redistribute rip route-map rip-routes
switch(config-ospf-1)# redistribute ospf 2
```

Disabling redistributing routes to OSPFv2:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no redistribute bgp
switch(config-ospf-1)# no redistribute bgp route-map BGP_routes
```

```
 switch(config-ospf-1)# no redistribute host-routes
switch(config-ospf-1)# no redistribute connected
switch(config-ospf-1)# no redistribute connected route-map connected_routes
switch(config-ospf-1)# no redistribute local loopback
```

```
switch(config-ospf-1)# no redistribute local loopback route-map local_routes
switch(config-ospf-1)# no redistribute static
switch(config-ospf-1)# no redistribute static route-map static_networks
switch(config-ospf-1)# no redistribute rip
switch(config-ospf-1)# no redistribute rip route-map rip-routes
switch(config-ospf-1)# no redistribute ospf 2
```

Redistributing ARP entries:

```
switch(config-ospf-1)# redistribute host-routes route-map host_routes12
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.13	Added **route-map** support for **host-routes**.
10.12	Added information about the **host-routes** parameter.
10.08	Added **route-map** support for supported redistribute source-protocols. Updated information and examples.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospf-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

reference-bandwidth

```
reference-bandwidth <BANDWIDTH>
no reference-bandwidth
```

Description

Sets the reference bandwidth for OSPFv2. If the OSPFv2 interface cost is not explicitly set, then the cost of all the OSPFv2 interfaces is recalculated based on the reference bandwidth and link speed of the interface.

For VLAN interfaces the link speed value is taken as 1 Gbps (if the OSPFv2 interface cost is not explicitly set).

The **no** form of this command sets the reference bandwidth for OSPF to the default of 100000 Mbps.

Parameter	Description
<BANDWIDTH>	Specifies the reference bandwidth used to calculate the cost of an interface in Mbps. Range: 1 to 4000000. Default: 100000.

Examples

Setting the reference bandwidth:

```
switch(config)# router ospf 1
switch(config-ospf-1)# reference-bandwidth 40000
```

Setting the reference bandwidth to the default value:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no reference-bandwidth
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospf-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

retransmit-interval

```
retransmit-interval <INTERVAL>
no retransmit-interval
```

Description

Sets the time between retransmitting lost link state advertisements for virtual links.

The **no** form of this command sets the time between retransmitting lost link state advertisements to the default of 5 seconds for virtual links.

Parameter	Description
<INTERVAL>	Specifies the retransmit interval in seconds. Range: 1 to 3600. Default: 5.

Examples

Setting OSPFv2 virtual links retransmit interval:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# retransmit-interval 30
```

Setting OSPFv2 virtual links retransmit interval to default:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# no retransmit-interval
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-router-vlink`	Administrators or local user group members with execution rights for this command.

rfc1583-compatibility

```
rfc1583-compatibility
no rfc1583-compatibility
```

Description

Enables OSPF compatibility with RFC1583 (backward compatibility). If RFC1583 compatibility is enabled, then the route cost calculation follows a different method.

The **no** form of this command disables OSPF compatibility with RFC1583 (backward compatibility). By default the RFC1583 compatibility is disabled.

Examples

Enabling OSPF RFC1583 compatibility:

```
switch(config)# router ospf 1
switch(config-ospf-1)# rfc1583-compatibility
```

Disabling OSPF RFC1583 compatibility:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no rfc1583-compatibility
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospf-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

router ospf

```
router ospf <PROCESS-ID> [vrf <VRF-NAME>]
no router ospf <PROCESS-ID> [vrf <VRF-NAME>]
```

Description

Creates an OSPF process (if not created already) on a VRF, and switches to the OSPF router instance context. Up to eight OSPF processes are supported per VRF.

The **no** form of this command removes the OSPF instance.

Parameter	Description
<PROCESS-ID>	Specifies an OSPF process ID. Range: 1 to 65535.

Parameter	Description
`vrf <VRF-NAME>`	Specifies a VRF name for the OSPF process. Default: default.

Examples

```
switch(config)# router ospf 1
switch(config-ospf-1)#
```

```
switch(config)# router ospf 1 vrf vrf_red
```

```
switch(config)# no router ospf 1
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

router-id

```
router-id <ROUTER-ADDR>
no router-id
```

Description

Sets an ID for the router in an IPv4 address format.

The **no** form of this command unconfigures the router-id for the instance and sets the router-id to the default as follows: the router-id is selected dynamically as equal to the highest loopback address on the router, or the highest active interface if there are no loopback addresses. If no IP address is configured on any interfaces on the router, OSPF will not form an adjacency.

Parameter	Description
<ROUTER-ADDR>	Specifies the Router ID in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.

Examples

Setting router-id in the OSPF context:

```
switch(config)# router ospf 1
switch(config-ospf-1)# router-id 1.1.1.1
```

Unconfiguring router-id:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no router-id
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospf-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

sha-key sha

```
sha-key <KEY-ID> sha [{ciphertext | plaintext} <KEY>]
no sha-key <KEY-ID>
```

Description

Sets the SHA (secure hash authentication) key for the selected virtual link. If the SHA key is given in ciphertext, it will be decrypted and applied to the protocol. This command accepts a key of up to 64 characters irrespective of the SHA version configured on virtual link. OSPF will internally pad zeros to the key to obtain a 64-byte key. For all types of SHA, key length is adjusted to 64 bytes.

The **no** form of this command deletes the virtual link SHA authentication key.

Parameter	Description
<KEY-ID>	Specifies the virtual link SHA key ID. Range: 1 to 255.
`{ciphertext | plaintext}`	Selects the virtual link SHA key format.
<KEY>	Specifies the virtual link SHA authentication key.

When the authentication key is not provided on the command line, plaintext key prompting occurs upon pressing Enter. The entered key characters are masked with asterisks.

Examples

Setting virtual link SHA authentication key in plaintext format:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# sha-key 1 sha plaintext F82#450b
```

Setting the virtual link SHA authentication key in prompted plaintext format:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# sha-key 1 sha
Enter the SHA authentication key: ********
Re-Enter the SHA authentication key: ********
```

Setting the virtual link SHA authentication key in ciphertext format:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# sha-key 1 sha ciphertext AQapu...C2K47A=
```

Deleting the virtual link SHA authentication key:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# no sha-key 1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-router-vlink`	Administrators or local user group members with execution rights for this command.

show ip ospf

`show ip ospf [<PROCESS-ID>] [all-vrfs | vrf <VRF-NAME>]`

Description

Displays general OSPF, area, state, and configuration information.

Parameter	Description
<PROCESS-ID>	Enter an OSPF process ID to display general OSPF information for a particular OSPF process. Range: 1 to 65535.
`all-vrfs`	Optionally select to display general OSPF information for all VRFs.
`vrf <VRF-NAME>`	Specifies the name of a VRF. Default: default.

Examples

Showing general OSPF configurations:

```
switch# show ip ospf 200
VRF : Default                           Process : 200
------------------------------------------------------------------------

Router ID              : 1.1.1.1            OSPFv2                 : Enabled
BFD                    : Disabled           SPF Start Interval     : 5000 ms
SPF Hold Interval      : 1000 ms            SPF Max Wait Interval  : 1000 ms
LSA Start Interval     : 5000 ms            LSA Hold Interval      : 1000 ms
LSA Max Wait Interval  : 1000 ms            LSA Arrival Interval   : 1000 ms
External LSAs          : 4                  Checksum Sum           : 133302
ECMP                   : 4                  Reference Bandwidth    : 100000 Mbps
Area Border            : Yes                AS Border              : No
GR Status              : Disabled           GR Interval            : 120 sec
GR State               : Inactive           GR Exit Status         : None
GR Helper              : Enabled            GR Strict LSA Check    : Enabled
GR Ignore Lost I/F     : Disabled           Internal Process ID    : 1
Summary address:
  prefix 10.1.1.0/24, advertise, tag 10

Area      Total    Active
-------------------------
```

```
Normal    2        1
Stub      2        1
NSSA      0        0

Area  : 0.0.0.1
----------------------------------
Area Type             : Normal     Status                 : Active
Total Interfaces      : 100        Active Interfaces      : 10
Passive Interfaces    : 5          Loopback Interfaces    : 85
SPF Calculation Count : 1500
Area ranges:
  ip-prefix 10.1.1.1/24, inter-area, advertise
Number of LSAs        : 5000       Checksum Sum           : 99122
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ip ospf border-routers

```
show ip ospf [<PROCESS-ID>]
     border-routers [all-vrfs | vrf <VRF-NAME>]
```

Description

Displays the OSPF routing table entries for Area Border Router (ABR) and Autonomous System Border Router (ASBR).

Parameter	Description
<PROCESS-ID>	Enter an OSPF process ID to display general OSPF information for a particular OSPF process. Range: 1 to 65535.
`all-vrfs`	Optionally select to display general OSPF information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Examples

Showing OSPF border routers information:

```
switch# show ip ospf border-routers
VRF : default                          Process : 1
Internal Routing Table
-----------------------------------------------------

Codes: i - Intra-area route, I - Inter-area route

  Router-ID    Cost Type Area      SPF   Nexthop    Interface
i 40.40.40.40  10   ABR  0.0.0.0   71    192.0.2.1  1/1/1
i 60.60.60.60  20   ABR  0.0.0.0   71    192.0.2.1  1/1/1
i 40.40.40.40  10   ABR  0.0.0.1   71    192.0.2.1  1/1/2
i 60.60.60.60  20   ABR  0.0.0.1   71    192.0.2.1  1/1/2
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ip ospf interface

```
show ip ospf [<PROCESS-ID>] interface [<interface-name>] [brief]
[all-vrfs | vrf <VRF-NAME>]  [vsx-peer]
```

Description

Displays general OSPF, area, state, and configuration information.

For an IP unnumbered enabled OSPF interface, the output of the command displays the IP address borrowed from the donor interface.

Parameter	Description
<PROCESS-ID>	Enter an OSPF process ID to display general OSPF information for a particular OSPF process. Range: 1 to 65535.
<interface-name>]	Specify the name of an OSPF interface.
brief	Include this parameter to display a brief overview of the following OSPF configuration information. ■ Interface: OSPF interface name. ■ Area: OSPF area ID. ■ Cost: The metric OSPF uses to judge a path's feasibility, calculated as (reference bandwidth / interface bandwidth). ■ State: Indicates if the interface is a designated router (**Dr**) or a backup designated router (**Backup-dr**). ■ Status: Indicates if the interface is **up** or **down**. ■ Flags: P - Passive A - Active.
all-vrfs	Optionally select to display general OSPF information for all VRFs.
vrf *<VRF-NAME>*	Specifies the name of a VRF. Default: default.
vsx-peer	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Examples

Showing general OSPF configuration settings for the default VRF:

```
switch (config-if)# show ip ospf 1
VRF : default                          Process  : 1
-----------------------------------------------------

RouterID               : 20.0.0.1           OSPFv2                 : Enabled
SPF Start Interval     : 200   ms
SPF Hold Interval      : 1000  ms           SPF Max Wait Interval  : 5000  ms
LSA Start Time         : 5000  ms           LSA Hold Time          : 0     ms
LSA Max Wait Time      : 0     ms           LSA Arrival            : 1000  ms
External LSAs          : 0                  Checksum Sum           : 0
ECMP                   : 4                  Reference Bandwidth    : 100000 Mbps
Area Border            : false              AS Border              : false
GR Status              : Enabled            GR Interval            : 120 sec
GR State               : inactive           GR Exit Status         : none
GR Helper              : Enabled            GR Strict LSA Check    : Enabled
GR Ignore Lost I/F     : Disabled
Summary address:

Area      Total     Active
------------------------------
Normal     1         1
Stub       0         0
NSSA       0         0

Area  : 0.0.0.0
----------------
Area Type              : Normal         Status               : Active
Total Interfaces       : 1              Active Interfaces    : 1
```

```
Passive Interfaces     : 0              Loopback Interfaces  : 0
SPF Calculation Count  : 4
Area ranges     :
Number of LSAs         : 3              Checksum Sum         : 82420
```

Showing OSPF configuration settings for all interfaces:

```
switch(config)# show ip ospf interface
Codes: DR - Designated router BDR - Backup designated router

Interface 1/1/1 is up, line protocol is up
-------------------------------------------
  VRF                 : default               Process             : 1
  IP Address          : 10.10.10.1/24         Area                : 0.0.0.0
  Status              : Up                    Network Type        : Broadcast
  Hello Interval      : 10    sec             Dead Interval       : 40    sec
  Transit Delay       : 1     sec             Retransmit Interval : 5     sec
  Link Speed          : 1000 Mbps
  Cost Configured     : 1                     Cost Calculated     : 1
  State/Type          : BDR                   Router Priority     : 1
  DR                  : 10.10.10.2            BDR                 : 10.10.10.1
  Link LSAs           : 0                     Checksum Sum        : 0
  Authentication      : Md5                   Passive             : Yes
Interface 1/1/2 is up, line protocol is up
-------------------------------------------
  VRF                 : default               Process             : 1
  IP Address          : 10.10.10.1/24         Area                : 0.0.0.0
  Status              : Up                    Network Type        : Broadcast
  Hello Interval      : 10    sec             Dead Interval       : 40    sec
  Transit Delay       : 1     sec             Retransmit Interval : 5     sec
  Link Speed          : 1000 Mbps
  Cost Configured     : 1                     Cost Calculated     : 1
  State/Type          : BDR                   Router Priority     : 1
  DR                  : 20.10.10.2            BDR                 : 20.10.10.1
  Link LSAs           : 0                     Checksum Sum        : 0
  Authentication      : Simple                Passive             : No
```

Displaying brief OSPF information

```
switch(config-if)# show ip ospf interface brief
VRF : default   Process : 1
==============================
Total Number of Interfaces: 2
Flags: P - Passive  A - Active
Interface     Area             IP Address/Mask     Cost  State    Status   Flags
-------------------------------------------------------------------------------
-------
1/1/1         0.0.0.0          10.10.10.1/24       40    DR       Up       P
1/1/2         255.255.255.255  200.200.200.123/24  4     Waiting  Up       A
```

Command History

Release	Modification
10.14.1000	, The output of the command displays the borrowed IP address for an IP unnumbered enabled OSPF interface.

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.09	Output of the **show ip ospf interface** command includes flags to indicate whether the interface is in passive or active mode.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ip ospf lsdb

```
show ip ospf [<process-id>] lsdb
   adv-router {<ROUTER-ID> | self}
   area <AREA-ID>
   lsid <link-state-id>
   all-vrfs |{vrf <VRF-NAME>}
      asbr-summary
      database-summary
      external
      lsid <LINK-STATE-ID>
      network
      nssa-external
      router
      summary
```

Description

Shows the OSPF link state database summary for different OSPF LSAs (Link State Advertisement). Use the parameters to get information for a particular LSA.

Parameter	Description
<PROCESS-ID>	Enter an OSPF process ID to display general OSPF information for a particular OSPF process. Range: 1 to 65535.
adv-router {*<ROUTER-ID>*|self}	Select to display link states for a particular advertising router. Specify either a Router ID of the advertising router or specify **self** to show self-originated link states.
area *<AREA-ID>*	Select to display information filtered for the specified area in one of the following formats.

Parameter	Description
	OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Value: 0 to 4294967295.
`lsid <LINK-STATE-ID>`	Select to display information filtered by link state identifier specified in IPv4 address format (A.B.C.D).
`all-vrfs|{vrf <vrf-name>`	Select **all-vrfs** to display general OSPF information for all VRFs, or use the **vrf <*VRF-NAME*>** option to display information for a specific VRF. Optionally select one of the following parameters to filter the link state database information.
`asbr-summary`	Show ASBR summary link states (LSA type 4).
`database-summary`	Select to display the count of each type of LSA and each area in the database. **NOTE:** The **database-summary** parameter does not support the **area <area-id>**, **lsid <link-state-id>** or **adv-router {<router-id>|self}** parameters.
`external`	Show external link states (LSA type 5). **NOTE:** The **external** parameter does not support the area **area <area-id>** parameter.
`network`	Show network LSAs (LSA type 2).
`nssa-external`	Show NSSA external link states (LSA type 7).
`router`	Show router LSAs (LSA type 1).
`summary`	Show network-summary link states (LSA type 3).

Examples

Showing OSPF link state database (LSDB) general information:

```
switch# show ip ospf lsdb
OSPF Router with ID (50.50.50.50) (Process ID 1 VRF default)
============================================================

Router Link State Advertisements (Area 0.0.0.0)
------------------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID         Link Count  Bits
------------------------------------------------------------------------------
40.40.40.40     930     0x80000004  0x2ea1      0            3           None
50.50.50.50     935     0x80000002  0x8b52      0            1           E
60.60.60.60     943     0x800003c5  0x9854      0            2           None

Network Link State Advertisements (Area 0.0.0.0)
------------------------------------------------------------------------
```

```
ADV Router      Age     Seq#        Checksum    LSID        Router Count
-------------------------------------------------------------------------
60.60.60.60     944     0x80000001  0x7179      1360007168  2
50.50.50.50     935     0x80000001  0x516a      19          1

Inter Area Prefix Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
-------------------------------------------------------------------
40.40.40.40     929     0x80000001  0x2498      131072      FEC0:3344::/32
50.50.50.50     928     0x80000001  0x5b2f      65536       111::/64

Inter Area Router Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Destination Router ID
-------------------------------------------------------------------------------
40.40.40.40     929     0x80000001  0x2498      1           33.33.33.33

AS External Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
-------------------------------------------------------------------
40.40.40.40     264     0x80000001  0x24cc4     1           10::/64
40.40.40.40     675     0x80000001  0x5b00f     2           11::/64

NSSA External Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
-------------------------------------------------------------------
3.3.3.3         264     0x80000001  0x24ac2     1           200::/64

Link-local Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Interface
-------------------------------------------------------------------
50.50.50.50     264     0x80000001  0x653c4     19          1/1/1

Intra Area Prefix Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------------
-------------
ADV Router      Age     Seq#        Checksum    LSID        Referenced LS Type
Referenced LSID
-------------------------------------------------------------------------------
-------------
50.50.50.50     263     0x80000001  0x1da34     1           0x2001              0
50.50.50.50     264     0x80000001  0x2a45d     1           0x2002              19
```

Showing ASBR summary link states:

```
switch# show ip ospf lsdb asbr-summary
OSPF Router with ID (2.2.2.1) (Process ID 1 VRF default)
=========================================================

ASBR Summary Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------------

LSID            ADV Router      Age     Seq#            Checksum
-----------------------------------------------------------------
209.165.201.3   60.60.60.60     944     0x80000001      0x7179
192.0.2.1       50.50.50.50     935     0x80000001      0x516a
```

Showing external link states:

```
switch# show ip ospf lsdb external
OSPF Router with ID (2.2.2.1) (Process ID 1 VRF default)
==========================================================

AS External Link State Advertisements
-------------------------------------

LSID            ADV Router      Age     Seq#            Checksum
-------------------------------------------------------------------
209.165.201.3   60.60.60.60     944     0x80000001      0x7179
192.0.2.1       50.50.50.50     935     0x80000001      0x516a
```

Showing database summary:

```
switch# show ip ospf lsdb database-summary
OSPF Router with ID (10.1.1.1) (Process ID 1 VRF default)
==========================================================

Area 0.0.0.0 database summary
-----------------------------

LSA Type             Count
---------------------------
Router               2
Network              1
Inter-area Summary   1
ASBR Summary         0
NSSA External        0
Subtotal             4

Process 1 database summary
--------------------------

LSA Type             Count
---------------------------
Router               2
Network              1
Inter-area Summary   1
ASBR Summary         0
NSSA External        0
AS External          0
Total                4
```

Showing router LSAs:

```
switch# show ip ospf lsdb router
OSPF Router with ID (2.2.2.1) (Process ID 1 VRF default)
==========================================================

Router Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------

LSID            ADV Router      Age         Seq#       Checksum Link Count
-------------------------------------------------------------------------
1.1.1.2         1.1.1.2         15          0x80000004 0xf526    1
2.2.2.1         2.2.2.1         14          0x80000005 0x6c5e    2
2.2.2.2         2.2.2.2         104         0x80000004 0xf51a    1
OSPF Router with ID (2.2.2.1) (Process ID 1 VRF default)
==========================================================
```

```
Router Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------

LSID            ADV Router      Age         Seq#       Checksum Link Count
-------------------------------------------------------------------------
1.1.1.2         1.1.1.2         15          0x80000004 0xf526    1
2.2.2.1         2.2.2.1         14          0x80000005 0x6c5e    2
2.2.2.2         2.2.2.2         104         0x80000004 0xf51a    1
```

Showing network LSAs:

```
switch# show ip ospf lsdb network
OSPF Router with ID (2.2.2.1) (Process ID 1 VRF default)
========================================================

Network Link State Advertisements (Area 0.0.0.0)
------------------------------------------------

LSID            ADV Router      Age         Seq#       Checksum
-----------------------------------------------------------------
1.1.1.2         1.1.1.2         141         0x80000001 0xc55e
2.2.2.2         2.2.2.2         230         0x80000001 0xa179
```

Showing network-summary link states:

```
switch# show ip ospf lsdb summary
OSPF Router with ID (2.2.2.1) (Process ID 1 VRF default)
========================================================

Inter-area Summary Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------------------

LSID            ADV Router      Age         Seq#       Checksum
-----------------------------------------------------------------
1.1.1.0         2.2.2.1         133         0x80000002 0xa089

Inter-area Summary Link State Advertisements (Area 0.0.0.1)
-----------------------------------------------------------

LSID            ADV Router      Age         Seq#       Checksum
-----------------------------------------------------------------
2.2.2.0         2.2.2.1         133         0x80000002 0x7caa
```

Showing NSSA external link states:

```
switch(config-ospf-1)# show ip ospf lsdb nssa-external
OSPF Router with ID (2.2.2.1) (Process ID 1 VRF default)
========================================================

NSSA External Link State Advertisements (Area 0.0.0.1)
------------------------------------------------------

LSID            ADV Router      Age         Seq#       Checksum
-----------------------------------------------------------------
8.8.8.0         1.1.1.2         162         0x80000003 0xc7b2
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ip ospf neighbors

```
show ip ospf [<PROCESS-ID>] neighbors [<NEIGHBOR-ID>]
     [interface <INTERFACE-NAME>] [detail | summary]
     [all-vrfs | vrf <VRF-NAME>]
```

Description

Displays information about OSPF neighbors.

Parameter	Description
<PROCESS-ID>	Enter an OSPF process ID to display OSPF neighbor information for the particular OSPF process. Range: 1 to 65535.
`neighbors` *<NEIGHBOR-ID>*	Select to display information about a particular neighbor, specified in IPv4 format (A.B.C.D).
`interface` *<INTERFACE-NAME>*	Select to display neighbor information only for the specified interface.
`detail`	Select to display detailed information for all the neighbors.
`summary`	Select to display summary information for the neighbors.
`all-vrfs`	Select to display neighbor information for all VRFs.
`vrf` *<VRF-NAME>*	Specify the name of a VRF. Default: default.

Examples

Showing OSPF neighbors information for the default VRF:

```
switch# show ip ospf neighbors
OSPF Process ID 1 VRF default
==============================

Total Number of Neighbors: 1

Neighbor ID      Priority  State             Nbr Address       Interface
-------------------------------------------------------------------------
2.2.2.2          1         FULL/DR           10.1.1.2           1/1/1
```

The **Nbr Address** displays the borrowed IP address for IP unnumbered interfaces.

Showing OSPF neighbors information for VRF red:

```
switch# show ip ospf neighbors vrf red
OSPF Process ID 1 VRF red
==========================

Total Number of Neighbors: 1

Neighbor ID      Priority  State             Nbr Address       Interface
-------------------------------------------------------------------------
1.1.1.1          1         FULL/BDR          10.1.1.1           1/1/1
```

Showing OSPF neighbors information for a specific neighbor:

```
switch# show ip ospf neighbors 2.2.2.2
switch# show ip ospf neighbors 2.2.2.2
VRF : default                          Process : 1
-------------------------------------------------------------
Router-Id      : 2.2.2.2               Area Id           : 0.0.0.0
Interface      : 1/1/1                 Address           : 10.10.10.2
State          : FULL                  Neighbor Priority : 1
Dead Timer Due : 00:00:36              Options           : 0x42
```

Showing OSPF neighbors information for a specific neighbor and interface:

```
switch# show ip ospf neighbors 2.2.2.2 interface 1/1/1
VRF : default                          Process : 1
-----------------------------------------------

Router-Id      : 2.2.2.2               Area Id           : 0.0.0.0
Interface      : 1/1/1                 Address           : 10.10.10.2
State          : FULL                  Neighbor Priority : 1
Dead Timer Due : 00:00:36              Options           : 0x42
```

Showing detail information for OSPF neighbors:

```
switch# show ip ospf neighbors detail
VRF : default                          Process : 1
-----------------------------------------------

Router-Id      : 2.2.2.2               Area Id           : 0.0.0.0
Interface      : 1/1/1                 Address           : 10.10.10.2
```

```
State          : FULL                 Neighbor Priority : 1
DR             : 10.10.10.2           BDR               : 10.10.10.1
Dead Timer Due : 00:00:38             Options           : 0x42
Retransmission Queue Length  : 0
Time Since Last State Change : 00h:11m:37s
```

Showing summary information for OSPF neighbors in the **default** VRF:

```
switch# show ip ospf neighbors summary
OSPF Process ID 1 VRF default, Neighbor Summary
================================================

Interface  Down Attempt Init TwoWay ExStart Exchange Loading Full  Total
-------------------------------------------------------------------------
1/1/1      0    0       0    0      0       0        0       1     1
1/1/2      0    0       0    0      0       0        0       0     0
Total      0    0       0    0      0       0        0       1     1
```

Command History

Release	Modification
10.14.1000	The output parameter, **Nbr Address** displays the borrowed IP address for IP unnumbered interfaces.
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ip ospf routes

```
show ip ospf [<PROCESS-ID>] routes
     [<IPV4-ADDR>/<MASK>] protected [all-vrfs | vrf <VRF-NAME|
```

Description

Displays OSPF routing information.

Parameter	Description
<IPV4-ADDR>	Specify an IP address in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.
<MASK>	Specifies the number of bits in the address mask in CIDR format (**x**), where **x** is a decimal number from 0 to 32.
<PROCESS-ID>	Enter an OSPF process ID to display OSPF neighbor information for the particular OSPF process. Range: 1 to 65535.
`all-vrfs`	Select to display neighbor information for all VRFs.
`vrf` *<VRF-NAME>*	Specify the name of a VRF. Default: default.
`protected`	Selects a protected route. Specifying an IPv4 address will select protected route for that prefix.

Examples

Showing OSPF routing table information:

```
switch# show ip ospf routes
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF default, Routing Table
--------------------------------------------

Total Number of Routes : 2

10.1.1.0/24        (i) area: 0.0.0.0
     directly attached to interface 1/1/1, cost 1 distance 110
20.1.1.0/24        (I)
     via 10.1.1.2 interface 1/1/1, cost 2 distance 110
```

Showing OSPF routing table information for a specific subnet:

```
switch# show ip ospf routes 10.1.1.0/2
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF default, Routing Table for prefixes 10.1.1.0/24
---------------------------------------------------------------------

Total Number of Routes : 1

10.1.1.0/24        (i) area: 0.0.0.0
     directly attached to interface 1/1/1, cost 1 distance 110
```

Showing OSPF routing table information for a specific protected subnet on the default VRF:

```
switch# show ip ospf routes 211.0.0.0/24 protected vrf default
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1, VRF default, Protected Routes for prefixes 211.0.0.0/24
```

```
--------------------------------------------------------------------------------

211.0.0.0/24       (i) area: 0.0.0.0
        via 110.110.110.1, interface vlan20, cost 122
        bkp 130.130.130.1, interface vlan22, cost 132


Total Number of Routes : 1
```

Showing OSPF routing table information for a specific protected subnet on all VRFs

```
switch# show ip ospf routes 211.0.0.0/24 protected all-vrfs
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1, VRF default, Protected Routes for prefixes 211.0.0.0/24
--------------------------------------------------------------------------------

211.0.0.0/24       (i) area: 0.0.0.0
        via 110.110.110.1, interface vlan20, cost 122
        bkp 130.130.130.1, interface vlan22, cost 132


Total Number of Routes : 1
```

Showing OSPF routing table information for a specific protected subnet:

```
switch# show ip ospf routes 211.0.0.0/24 protected
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1, VRF default, Protected Routes for prefixes 211.0.0.0/24
--------------------------------------------------------------------------------

211.0.0.0/24       (i) area: 0.0.0.0
        via 110.110.110.1, interface vlan20, cost 122
        bkp 130.130.130.1, interface vlan22, cost 132

Total Number of Routes : 1
```

Showing OSPF routing table information for all protected routes:

```
switch# show ip ospf routes 211.0.0.0/24 protected
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1, VRF default, Protected Routes
--------------------------------------------------------------------------------

120.120.120.0/24       (i) area: 0.0.0.0
        via 110.110.110.1, interface vlan20, cost 120
        bkp 130.130.130.1, interface vlan22, cost 230

140.140.140.0/24       (i) area: 0.0.0.0
        via 110.110.110.1, interface vlan20, cost 130
        bkp 130.130.130.1, interface vlan22, cost 220
```

```
211.0.0.0/24       (i) area: 0.0.0.0
        via 110.110.110.1, interface vlan20, cost 122
        bkp 130.130.130.1, interface vlan22, cost 132

Total Number of Routes : 3

OSPFs Process ID 1, VRF red, Protected Routes
---------------------------------------------

Total Number of Protected Routes: 0
```

Showing OSPF routing table information for protected routes on the default VRF:

```
switch# show ip ospf routes protected vrf default
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1, VRF default, Protected Routes
------------------------------------------------------------------------------

140.140.140.0/24       (i) area: 0.0.0.0
        via 130.130.130.1, interface vlan22, cost 130
        bkp 110.110.110.1, interface vlan20, cost 220

211.0.0.0/24      (i) area: 0.0.0.0
        via 110.110.110.1, interface vlan20, cost 122
        bkp 130.130.130.1, interface vlan22, cost 132

120.120.120.0/24       (i) area: 0.0.0.0
        via 110.110.110.1, interface vlan20, cost 120
        bkp 130.130.130.1, interface vlan22, cost 230

Total Number of Routes : 3
```

Command History

Release	Modification
10.15	Introduces the **protected** parameter.
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P	Operator (>) or Manager (#)	

Platforms	Command context	Authority
8360 9300 9300S 10000		

show ip ospf statistics

`show ip ospf [<PROCESS-ID>] statistics [all-vrfs | vrf <VRF-NAME>]`

Description

Displays OSPF statistics.

Parameter	Description
<PROCESS-ID>	Enter an OSPF process ID to display OSPF neighbor information for the particular OSPF process. Range: 1 to 65535.
`all-vrfs`	Select to display OSPF statistics information for all VRFs.
`vrf <VRF-NAME>`	Specify the name of a VRF. Default: default.

Examples

Showing OSPF statistics:

```
switch# show ip ospf statistics
OSPF Process ID 1 VRF default, Statistics (cleared 1h 16m 24s ago)
------------------------------------------------------------

Unknown Interface Drops         : 0
Unknown Virtual Interface Drops : 0
Bad Instance ID Drops           : 0
Bad IP Header Length Drops      : 0
Wrong OSPF Version Drops        : 0
Bad Source IP Drops             : 0
Resource Failure Drops          : 0
Bad Header Length Drops         : 0
Total Drops                     : 0
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A,	Operator (>) or	

Platforms	Command context	Authority
S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Manager (#)	

show ip ospf statistics interface

```
show ip ospf [<PROCESS-ID>] statistics interface [<INTERFACE-NAME>]
     [all-vrfs | vrf <VRF-NAME>]
```

Description

Displays OSPF statistics for the OSPF-enabled interfaces.

Parameter	Description
<PROCESS-ID>	Enter an OSPF process ID to display OSPF-enabled interface statistics information on the specified OSPF process. Range: 1 to 65535.
<INTERFACE-NAME>	Select to display information only for the specified interface.
`all-vrfs`	Select to display OSPF-enabled interface statistics information for all VRFs.
`vrf <VRF-NAME>`	Specify the name of a VRF. Default: default.

Examples

Showing OSPF-enabled interfaces information:

```
switch# show ip ospf statistics interface 1/1/1
OSPF Process ID 1 VRF default, interface 1/1/1 statistics (cleared 0h 30m 28s ago)
================================================================================

Tx Hello Packets      : 101              Rx Hello Packets       : 99
Tx Hello Bytes        : 101              Rx Hello Bytes         : 99
Tx DD Packets         : 101              Rx DD Packets          : 99
Tx DD Bytes           : 101              Rx DD Bytes            : 99
Tx LS Request Packets : 101              Rx LS Requests Packets : 99
Tx LS Request Bytes   : 101              Rx LS Request Bytes    : 99
Tx LS Update Packets  : 101              Rx LS Update Packets   : 99
Tx LS Update Bytes    : 101              Rx LS Update Bytes     : 99
Tx LS Ack Packets     : 101              Rx LS Ack Packets      : 99
Tx LS Ack Bytes       : 101              Rx LS Ack Bytes        : 99

Total Number of State Changes : 8
Number of LSAs                : 29
LSA Checksum Sum              : 2345
Total Transmit Failures       : 29
```

```
Total OSPF Packets Discarded  : 999

Reason                                  Packets Dropped
-----------------------------------------------------
Invalid type                            19
Invalid length                          9
Invalid checksum                        0
Invalid version                         23
Bad or unknown source                   67
Area mismatch                           1
Self-originated                         19
Duplicate router ID                     9
Interface standby                       0
Total Hello packets dropped             60
  Network Mask mismatch                 10
  Hello interval mismatch               10
  Dead interval mismatch                10
  Options mismatch                      10
  MTU mismatch                          10
  Neighbor ignored                      10
Authentication errors                   12
  Type mismatch                         6
  Authentication failures               6
Wrong protocol                          0
Resource failures                       0
Bad LSA length                          0
Others                                  0

Total LSAs Ignored : 176
Bad Type           : 10
Bad Length         : 56
Invalid Data       : 55
Invalid Checksum   : 55
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ip ospf virtual-links

```
show ip ospf [<PROCESS-ID>] virtual-links [brief] [all-vrfs | vrf <VRF-NAME>]
```

Description

Displays the current state and parameters of the OSPF virtual links.

Parameter	Description
<PROCESS-ID>	Enter an OSPF process ID to display information on the OSPF virtual links for the particular OSPF process. Range: 1 to 65535.
`brief`	Select to display brief overview information for the OSPF virtual links.
`all-vrfs`	Select to display OSPF virtual links information for all VRFs.
`vrf <VRF-NAME>`	Specify the name of a VRF. Default: default.

Examples

Showing OSPF virtual links information:

```
switch# show ip ospf virtual-links
Virtual link to router 40.40.40.40 is up
----------------------------------------

VRF                 : default                       Process             : 21
Transit Area        : 0.0.0.1                       Authentication      : No
Hello Interval      : 10                            Dead Interval       : 40
Transit Delay       : 1                             Retransmit Interval : 5
Number of Link LSAs : 0                             Checksum Sum        : 0
Number of State Shanges : 4
```

Showing brief overview information for OSPF virtual links:

```
switch# show ip ospf virtual-links brief
OSPF Process ID 1 VRF default
==============================

Total Number of Virtual Links: 1

Remote Router    Transit Area      Status
----------------------------------------
2.2.2.2          0.0.0.1           down
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

summary-address

```
summary-address  <IPV4-ADDR>/<MASK> [no-advertise | tag <TAG-VALUE>]
no summary-address <prefix/length> [no-advertise | tag <tag-value>
```

Description

Summarizes the external routes with the matching address and mask. When advertising this route, its metric is set to the lowest cost path from among the routes that were summarized.

The **no** form of this command disables route summarization.

This command only works for an ASBR (Autonomous System Boundary Router).

Parameter	Description
<IPV4-ADDR>	Specifies an IP address in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.
<MASK>	Specifies the number of bits in the address mask in CIDR format (**x**), where **x** is a decimal number from 0 to 32.
`no-advertise`	Do not advertise the aggregate route. Suppress routes that match the specified prefix/mask pair.
`tag` *<TAG-VALUE>*	Specify the tag for the aggregate route. The summary prefix will be advertised along with the tag value in External LSAs. Range: 0 to 4294967295

Examples

Setting OSPF route summarization:

```
switch(config)# router ospf 1
switch(config-ospf-1)# summary-address 10.1.0.0/16
```

Disabling route summarization:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no summary-address 10.1.0.0/16
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospf-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

timers lsa-arrival

```
timers lsa-arrival <DELAY>
no timers lsa-arrival
```

Description

Configures the minimum delay between receiving the same LSA from a peer. The same LSA is an LSA that contains the same LSA ID number, LSA type, and advertising router ID. If an instance of the same LSA arrives sooner before the delay expires, the LSA is dropped. Generally, the LSA arrival timer should be set to a value less than or equal to the start-time value for the command **timers throttle lsa start** on the neighbor.

The **no** form of this command sets the LSA timers to default values.

Parameter	Description
<DELAY>	Specifies the delay in milliseconds. Range: 0 to 600000. Default: 1000.

Examples

Setting the LSA arrival timer:

```
switch(config)# router ospf 1
switch(config-ospf-1)# timers lsa-arrival 10
```

Setting the LSA arrival timer to default:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no timers lsa-arrival
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospf-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

timers throttle lsa

```
timers throttle lsa start-time <START-TIME> hold-time <HOLD-TIME> max-wait-time <WAIT-TIME>
no timers throttle lsa
```

Description

Configures the timers for LSA generation.

The **no** form of this command sets the LSA timers to default values.

Parameter	Description
`start-time <START-TIME>`	Specifies the initial wait time in milliseconds after which LSAs are generated. When set to 0, the LSAs are generated without any delay. Range: 0 to 600000. Default: 5000.
`hold-time <HOLD-TIME>`	Specifies the amount of time, in milliseconds, between regeneration of an LSA. The hold time doubles each time the same LSA must be regenerated, until **max-wait-time** is reached. When set to 0, LSA regeneration time is not increased. Range: 0 to 600000. Default: 0.
`max-wait-time <WAIT-TIME>`	Specifies the maximum wait time, in milliseconds, for regeneration of the same LSA. When set to 0, LSA regeneration time is not increased. Range: 0 to 600000. Default: 0.

Examples

Setting the LSA timers:

```
switch(config)# router ospf 1
switch(config-ospf-1)# timers throttle lsa start-time 100 hold-time 1000 max-wait-time 10000
```

Setting LSA timers to default values:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no timers throttle lsa
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospf-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

timers throttle spf

```
timers throttle spf start-time <milliseconds> hold-time <milliseconds>
max-wait-time <milliseconds>
no timers throttle spf
```

Description

Configures timers for SPF calculation. There are three timers:

- **start-time**: Is the initial delay before an SPF calculation is started. Default is 200 milliseconds.
- **hold-time**: Is the progressive backoff time to wait before next scheduled SPF calculation. Default is 1000 milliseconds. If a route change event occurs during this period, the value doubles until it reaches the *max-wait-time*.
- **max-wait-time:** Is the maximum time to wait before the next scheduled SPF calculation. Default is 5000 milliseconds. This is used to limit the SPF hold timer and also defines the time to be considered for which the OSPF LSDB has to be stable, after which the SPF throttle mechanism is reset.

The **no** form of this command sets all the configured non-default timers to default value.

Parameter	Description
<START-TIME>	Time in milliseconds to set timer for initial SPF delay. Default: 200. Range: 1-600000

Parameter	Description
	Required
<HOLD-TINME>	Time in milliseconds to set the minimum hold time between two consecutive SPF calculations. Default: 1000. Range: 1-600000 *Required*
<WAIT-TIME>	Time in milliseconds to set the maximum wait time between two consecutive SPF calculations. Default: 5000. Range: 1-600000 *Required*

Examples

Setting non-default timer values for SPF throttling:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# timers throttle spf start-time 100 hold-time 1000 max-
wait-time 10000
```

Setting default timer values for SPF throttling after configuring non-default values:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# no throttle spf
switch(config-ospfv3-af-1)# no timers throttle spf start-time 100 hold-time 1000
max-wait-time 10000
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-af-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

transit-delay

```
transit-delay <SECONDS>
no transit-delay
```

Description

Sets the time delay in Link state transmission for virtual links.

The **no** form of this command sets the delay in Link state transmission to the default of 1 second for virtual links.

Parameter	Description
<SECONDS>	Specifies the time delay for the transit delay, in seconds. Default: 1 second. Range: 1-3600.

Examples

Setting OSPFv2 virtual links transit delay:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# transit-delay 30
```

Setting OSPFv2 virtual links transit delay to default:

```
switch(config)# router ospf 1
switch(config-ospf-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# no transit-delay
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-router-vlink	Administrators or local user group members with execution rights for this command.

trap-enable

```
trap-enable
no trap-enable
```

Description

Enables the notification of the events to be sent as traps to the SNMP management stations for OSPF.

The **no** form of this command disables the notification of the events to be sent as traps to the SNMP management stations for OSPF.

Examples

Enabling sending notification of events as traps:

```
switch(config)# router ospf 1
switch(config-ospf-1)# trap-enable
```

Disabling sending notification of events as traps:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no trap-enable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospf-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

Chapter 6

Open Shortest Path First version 3

OSPFv3 is the IPv6 implementation of Open Shortest Path First protocol (OSPFv2 is the IPv4 implementation of this protocol). OSPFv3 is a routing protocol which is described in RFC5340 entitled OSPF for IPv6. It is a link-state based IGP (Interior Gateway Protocol) routing protocol. It is widely used with medium to large-sized enterprise networks.

Overview

The characteristics of OSPFv3 are:

- Provides a loop-free topology using SPF algorithm.
- Allows hierarchical routing using area 0 (backbone area) as the top of the hierarchy.
- Supports load balancing with equal cost routes for same destination.
- OSPFv3 is a classless protocol and allows for a hierarchical design with VLSM (Variable Length Subnet Masking) and route summarization.
- Scales enterprise size network easily with area concept.
- Provides fast convergence with triggered, incremental updates through Link State Advertisements (LSAs).

Some OSPFv3 configuration is done in the configuration context, others in the OSPFv3 router context, or in the interface context. OSPFv3 can be configured on routed ports, VLAN interfaces, LAG interfaces, and loopback interfaces. All such configurations work in the mentioned interfaces context. OSPFv3 mandates the associated interface to be routing interface.

The switch supports congestion control (prioritizing hello packets, inactivity timer reset, and adjacency throttling).

Supported features

- OSPFv3 neighbor adjacency, Hello protocol, multiple areas, Inter-area routing
- AS external routes, stub areas, totally stubby areas, NSSA, ABR, ASBR
- Designated Router/Backup Designated Router
- Point-to-point interfaces/broadcast interfaces
- Virtual Links
- Bidirectional Forwarding Detection (BFD) - refer to the *High Availability Guide* for additional information
- Equal-cost multipath
- Null authentication, Simple password authentication, MD5 authentication
- Area range aggregation - Type-3/Type-7
- External route aggregation - Type-5 LSAs
- Graceful Restart - un-planned, restart interval, Graceful Restart Helper
- Stub router advertisement
- SPF throttling

- LSA Throttling
- SPF throttling
- Graceful Restart - un-planned, restart interval, Graceful Restart Helper
- Stub router advertisement
- Configuration of interface parameters such as priority, cost, hello-interval, dead-interval, retransmit-interval, transit-delay, etc.
- Configuration of virtual link parameters such as hello-interval, dead-interval, retransmit-interval, transit-delay, etc.
- Route redistribution
- Passive interfaces
- Multi VRF support with each VRF having up to eight OSPF process instances
- Congestion control (prioritizing hello packets, inactivity timer reset, and adjacency throttling)

An OSPFv3 session will not come up in 6in6 tunnel.

How OSPFv3 protocol works

OSPFv3 protocol

The protocol uses Link State Advertisements (LSAs) transmitted by each router to update neighboring routers regarding its interfaces and the routes available through those interfaces. Each routing switch in an area also maintains a link-state database (LSDB) that describes the area topology. (All routers in a given OSPF area have identical LSDBs.) The routing switches used to connect areas to each other flood summary link LSAs and external link LSAs to neighboring OSPF areas to update them regarding available routes. Through this means, each OSPF router determines the shortest path between itself and a desired destination router in the same OSPF domain (Autonomous System (AS)).

OSPFv3 concepts

OSPFv3 is a link-state routing protocol applied to IPv6 routers grouped into OSPFv3 areas identified by the IPv6 routing configuration on each routing switch. Each OSPFv3 area includes one or more networks. OSPFv3 routers use hello packets and LSAs to maintain OSPFv3 operation across networks within an area and between areas within an OSPFv3 domain.

OSPFv3 uses hello packets to initiate and preserve relationships between neighboring routers on the same interface.

OSPFv3 Link-state advertisement (LSA) types

OSPFv3 uses LSAs transmitted by each router to update neighboring routers regarding its interfaces and the routes available through those interfaces.

Each routing switch in an area also maintains an LSDB that describes the area topology. (All routers in a given OSPFv3 area have identical LSDBs, and each router uses the LSDB to build its own shortest-path tree.)

The routing switches used to connect areas to each other flood inter-area-prefix-LSAs, inter-area-router-LSAs, and AS-external-LSAs to backbone area to update it regarding available routes. Each OSPFv3 router determines the shortest path between itself and a desired destination router in the same OSPFv3 domain (AS). Routed traffic in an OSPFv3 AS is classified as one of the following:

- Intra-area traffic
- Inter-area traffic
- External traffic

The routing switches support the LSAs listed in the following table.

Link-state type	Description	Use	Flood scope
0x2001	Router-LSA	Describes the state of each active interface on a router for Area a given area. (Excludes loopback interfaces and interfaces that have not achieved full adjacency.)	Area
0x2002	Network-LSA	Describes the OSPFv3 routers in a given network.	Area
0x2003	Inter-area-prefix-LSA	Describes the route to a prefix in another OSPFv3 area of Area the same AS. (Excludes prefixes for link-local addresses.) Propagated through backbone area to other areas.	Area
0x2004	Inter-area-router-LSA	Describes the route to an ASBR in another OSPFv3 normal area (including the backbone area) of the same AS. Propagated through backbone area to other areas. (Excludes any ASBR in the same area as the router sending the LSA.)	Area
0x4005	AS-external-LSA	Describes the route to a destination prefix in another AS (external route). (Excludes prefixes for link-local addresses.) Originated by ASBR in normal or backbone areas of an AS and propagates through backbone area to other normal areas. Does not flood over virtual links and is not summarized in virtual links. For injection into an NSSA, an NSSA ABR generates a type-7-default-LSA advertising the default route (::/0).	AS

Link-state type	Description	Use	Flood scope
0x2007	NSSA-LSA	Describes the route to a destination in another AS (external route). Originated by ASBR in NSSA. ABR translates type-7 LSAs to AS-external-LSAs for injection into the backbone area.	NSSA
0x0008	Link-LSA	For other routers on the same VLAN interface (except virtual links), describes the router link-local address and any other IPv6 prefixes reachable on the VLAN. Link LSAs are not flooded over virtual links.	Link-local
0x2009	Intra-area-prefix-LSA	Generated on transit links within an area by the DR operating on those links. Also, every OSPFv3 router generates this LSA to refer to stub and loopback prefixes on the router.	Area

OSPFv3 area types

OSPFv3 is built upon a hierarchy of network areas. All areas for a given OSPFv3 domain reside in the same AS. An AS is defined as a number of contiguous networks that share an interior gateway routing protocol.

An AS can be divided into multiple areas, including the backbone (area 0). Because each area represents a collection of contiguous networks and hosts, the topology of a given area is not known by the internal routers in any other area. Areas define the boundaries to which router-LSAs and network-LSAs are broadcast. This functionality limits the amount of LSA flooding that occurs within the AS. It also helps to control the size of the link-state databases (LSDBs) maintained in OSPFv3 routers.

An area is represented in OSPFv3 by either a 32-bit dotted-decimal address or a number. Area types include: Backbone, Normal, Not-so-stubby (NSSA), and Stub.

Backbone area

Every AS must have one (and only one) backbone area (identified as area 0 or 0.0.0.0.) The ABRs of all other areas in the same AS connect to the backbone area, either physically through an ABR or through a configured, virtual link. The backbone is a special type of area and serves as a transit area for carrying the inter-area-prefix-LSAs, AS-external-LSAs, and routed traffic between non-backbone areas, as well as the router-LSAs, network-LSAs, and routed traffic internal to the area. ASBRs are allowed in backbone areas.

Normal area

A normal area allows inter-area-prefix-LSAs and AS-external-LSAs to and from the backbone area. A normal area connects to the AS backbone area through one or more ABRs (physically or through a

virtual link) and allows router-LSAs and network LSAs within this area. ASBRs are allowed in normal areas.

Stub area

Stub area connects to the AS backbone through one or more ABRs. It does not allow an internal ASBR, and does not allow AS-external-LSAs. A stub area supports these actions:

- Advertise the area summary routes to the backbone area.
- Advertise inter-area routes from other areas.
- Use the inter-area-prefix-LSA default route to advertise routes to an ASBR and to other areas.

You can configure the stub area ABR to do the following:

- Suppress advertising on the area's summarized internal routes into the backbone area.
- Suppress LSA traffic from other areas in the AS by replacing inter-area-prefix-LSAs and the default external route from the backbone area with the default summary route (::/0.). This area is called totally stubby area.

Virtual links are not allowed for stub areas.

Not-so-stubby (NSSA) area

An NSSA area connects to the backbone area through one or more ABRs. NSSAs are used where an ASBR exists in an area where you want to:

- Block injection of external routes from other areas of the AS.
- Advertise type-7-LSA external routes (learned from the ASBR) to the backbone area as AS-external-LSAs.

NSSAs also support the following:

- Advertise inter-area-prefix-LSAs from the backbone area into the NSSA. (If no-summary is enabled, the NSSA ABR suppresses these LSAs from the backbone and, instead, injects the default route into the NSSA.)
- Advertise NSSA inter-area-prefix-LSAs to the backbone area.

Virtual links are not allowed for NSSAs.

OSPFv3 router types

Internal routers

Internal OSPFv3 routers belong to only one area. Internal routers flood router-LSAs to all routers in the same area and maintain identical LSDBs.

Area border routers (ABRs)

Area border routers have membership in multiple areas. ABRs are used to connect the various areas in an AS to the backbone area for that AS. Multiple ABRs can be used to connect a given area to the backbone, and a given ABR can belong to multiple areas other than the backbone.

An ABR maintains a separate LSDB for each area to which it belongs. (All routers within the same area have identical LSDBs.) The ABR is responsible for flooding inter-area-prefix-LSAs and inter-area router LSAs between its border areas.

You can reduce summary LSA flooding by configuring area ranges. An area range enables you to assign an aggregate address to a range of IPv6 addresses. This aggregate address is advertised instead of all the individual addresses it represents. You can assign up to eight ranges in an OSPFv3 area.

Autonomous system boundary router (ASBR)

Autonomous system boundary routers run one or more interior gateway protocols and serve as a gateway to other autonomous systems operating with interior gateway protocols. The ASBR imports and translates different protocol routes into OSPFv3 through redistribution. ASBRs can be used in backbone areas, normal areas, and NSSAs, but not in stub areas.

Designated routers (DRs)

In an OSPFv3 network having two or more routers, one router is elected to serve as the designated router (DR) and another router to act as the backup designated router (BDR). All other routers in the same network segment forward their routing information to the DR and BDR, and the DR forwards this information to all routers in the network. This functionality minimizes the amount of repetitive information that is forwarded on the network by eliminating the need for each individual router in the area to forward its routing information to all other routers in the network. If the area includes multiple networks, each network elects its own DR and BDR.

In an OSPFv3 network with no DR and no BDR, the neighboring router with the highest priority is elected the DR and the router with the next highest priority is elected the BDR. If the DR goes off-line, the BDR automatically becomes the DR, and the router with the next highest priority then becomes the new BDR. If multiple routing switches on the same OSPFv3 network are declaring themselves DRs, both priority and router ID are used to select the DR and BDR.

Priority is configurable using the `ipv6 ospfv3 priority` command at the interface level. You can use this command to help bias one router as the DR. If two neighbors share the same priority, the router with the highest router ID is designated the DR. The router with the next highest router ID is designated the BDR. Routers retain DR/BDR states throughout the period for which the adjacency is up and running.

OSPFv3 configuration task list

Tasks at a glance

- Configuring OSPFv3 on the routing switch
- Assigning the routing switch to an OSPF area
- Setting OSPFv3 network for the area
- Configuring external route redistribution and control
- Configuring area ranges on an ABR to reduce advertisements to the backbone
- Influencing route choice by changing the administrative distance
- Configuring graceful restart
- Configuring OSPFv3 virtual link settings
- Configuring OSPFv3 interface settings
- Configuring BFD for OSPFv3
- Configuring all OSPFv3 interfaces as passive
- Viewing OSPFv3 information
- Clearing OSPFv3 statistics on a switch

Configuring OSPFv3 on the routing switch

Prerequisites

- You must be in the global configuration context, as indicated by the `switch(config)#` prompt to create the OSPFv3 instance and enter the OSPFv3 configuration context.
- To configure a router ID, and other OSPFv3 configuration you must be in the OSPFv3 router configuration context, as indicated by the `switch(config-ospfv3-1 )#` prompt.

Create the OSPFv3 instance and enter the OSPFv3 configuration context. From this context, you can proceed with other OSPFv3 configuration.

Procedure

1. Create the OSPFv3 instance and enter the OSPFv3 configuration context using the following command. For command details, see router ospfv3.
   `router ospfv3 <PROCESS-ID> [vrf <VRF-NAME>]`
   For example, the following command creates OSPF instance 1.

   ```
   switch(config)# router ospfv3 1
   switch(config-ospfv3-1)#
   ```

2. Configure a global router ID using the following command. For command details, see router-id .
   `router-id <ROUTER-ADDRESSW>`
   For example, the following command sets the router ID to 1.1.1.1.

   ```
   switch(config-ospfv3-1)# router-id 1.1.1.1
   ```

3. Optionally, if the OSPFv3 process was disabled (it is enabled by default), enable the OSPFv3 process using the following command.
   `enable`
   For command details, see enable . (Refer to disable for disabling the OSPFv3 process).

Creating an OSPFv3 area

Prerequisites

You must be in the OSPFv3 router configuration context.

Create a Normal, Stub, or Not So Stubby (NSSA) area.

Procedure

Create an OSPFv3 area for the routing switch using one of the following commands:

- Create a Normal area using the following command: `area <area-id>`. For command details, see area .
- Create a Stub area using the following command: `area <area-id> stub`. For command details, see area stub .
- Create a Not So Stubby Area using the command: `area <area-id> nssa`.

For example, the following command creates a normal area with an area identifier of 10.1.1.1. Area identifier could alternatively be entered in decimal format such as 1.

```
switch(config-ospfv3-1)# area 1.1.1.1
```

Section	Owner-Group
Important Information	NTL/Dev
Products Supported (Define New HW Platforms)	PLM
Compatibility/interoperability	NTL/Dev
Transceiver Support	PLM
Enhancements and new features	PLM
Upgrade	NTL
Fixes (Defects)	
Define CFDs to include	PMO
Define IFDs to include	NTL
Issues & Workarounds	
Define IFDs to include	NTL

Setting OSPFv3 network for the area

Prerequisites

Routing must be enabled on the interface where you are planning to enable OSPFv3.

You must be in the interface configuration context, as indicated by the `switch(config-if)#` prompt.

After you define an OSPFv3 area, you can assign one or more networks to it. The OSPFv3 protocol will run on the interface with the configured link-local ipv6 address. The interfaces that have an IPv6 address configured in this network or in a subset of this network will participate in the OSPFv3 protocol.

Procedure

1. Set an OSPFv3 network for the area using the following command. For command details, see ipv6 ospfv3 area.
   `ipv6 ospfv3 <PROCESS-ID> area <AREA-ID>`
   For example, use the following command to assign interface 1/1/1 to OSPFv3 area 1. The area can alternatively be entered in IO address format.

```
switch(config)# interface 1/1/1
switch(config-if)# ipv6 ospfv3 1 area 1
```

2. Optionally, you can disable OSPFv3 on the interface using the following command. For command details, see ipv6 ospfv3 area.
   `ipv6 ospfv3 shutdown`

```
switch(config)# interface 1/1/1
switch(config-if)# ipv6 ospfv3 shutdown
```

Configuring external route redistribution and control

Configuring route redistribution for OSPFv3 establishes the routing switch as an ASBR for importing and translating different protocol routes from other IGP domains into an OSPFv3 domain. When you configure redistribution for OSPFv3, you can specify that routes external to the OSPFv3 domain are imported as OSPFv3 routes.

1. Enable route redistribution using the `redistribute` command.
   `redistribute {bgp | connected | local loopback | static | ripng | ospf <PROCESS-ID>}`
   `[route-map <ROUTE-MAP-NAME>]`

   For example, setting redistribution of connected routes as OSPFv3 routes:

```
switch(config-ospfv3-1)# redistribute connected
```

   If a route map is specified, then only the routes that pass the match clause specified in the route map are redistributed to OSPFv3. For example, setting redistribution of local loopback routes, that only match the routes specified by the route map:

```
switch(config-ospf-1)# redistribute local loopback route-map local_routes
```

2. Optionally, modify the default metric for redistribution using the `default-metric` command.
   `default-metric <METRIC-VALUE>`

   For example, setting the default metric for redistribution to 37.

```
switch(config-ospfv3-1)# default-metric 37
```

3. Optionally, set the cost of default-summary LSAs using the area command.
   `area <AREA-ID> default-metric <COST>`

   For example, setting the cost of default summary LSAs to 2.

```
switch (config-ospfv3-1)# area 1 default-metric 2
```

4. Optionally, use the `max-metric router-lsa` command to set the protocol to advertise a maximum metric so that other routers do not prefer this router as an intermediate hop in their shortest path first (SPF) calculations.
   `max-metric router-lsa [on-startup [<ADVERT-TIME>]]`

For example, setting advertise max-metric router-lsa on startup.

```
switch(config-ospfv3-1)# max-metric router-lsa on-startup 3000
```

5. Optionally set the maximum number of ECMP routes that OSPFv3 can support using the `maximum-paths` command.
   `maximum-paths <MAX-VALUE>`
   For example, setting the maximum number of ECMP routes to 8.

   ```
   switch(config-ospfv3-1)# maximum-paths 8
   ```

Configuring area ranges on an ABR to reduce advertisements to the backbone

You can configure area ranges to reduce inter-area advertisements by summarizing a range of IP addresses into a single route advertisement. This action prevents an ABR from advertising specific networks or subnets to the backbone area.

Prerequisites

You must be in the router configuration context, as indicated by the `switch(config-ospfv3)#` prompt.

Procedure

Summarize inter-area or NSSA paths using the following command. For command details, see area range .

`area <area-id> range <ip-prefix> type {inter-area | nssa} [no-advertise]`

For example, use the following command to summarize routes matching the area range 172.77.114.0/24 using inter-area as the type of address aggregation.

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 1 range fd00::/64 type inter-area
```

In another example, use the following command to specify DoNotAdvertise status for routes matching the area range fd00::/64. Use nssa as the type of address aggregation.

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 1 range fd00::/64 type inter-area no-advertise
```

Influencing route choice by changing the administrative distance

Prerequisites

You must be in the OSPFv3 router configuration context, as indicated by the `switch(config-ospfv3-1)#` prompt.

The administrative distance is used to influence the selection of routes learned by different protocols.

Procedure

Reconfigure the administrative distance using the following command. For command details, see distance .

`distance <distance>`

For example, use the following command to set administrative distance to 100.

```
switch(config-ospfv3-1)# distance 100
```

Configuring graceful restart

Prerequisites

Graceful restart is only applicable to the 6300-VSF and 6400, not for 832x switches.

You must be in the OSPFv3 router configuration context, as indicated by the `switch(config-ospfv3-1)#` prompt.

OSPFv3 routing can be gracefully restarted on switches without losing packets that are in transit. There is no effect on the saved switch configuration.

Manually create a router-id configuration to support graceful restart and management module failover. A dynamic router-id is chosen from the list of configured IPs at the time of the router-id election, and in the event of a graceful restart and management module failover, the router-id should be preserved before and after these events. Graceful restart would fail if there is a change in the router-id, causing traffic loss.

Procedure

Configure graceful restart of using the following command. For command details, see graceful-restart .

`graceful-restart [restart-interval <seconds> | helper]`

For example, the following command specifies 50 seconds as the maximum interval another router will wait for this router to gracefully restart.

```
switch(config-ospfv3-1)# graceful-restart restart-interval 50
```

Configuring OSPFv3 virtual link settings

You must be in the router vlink configuration context, as indicated by the `switch(config-router-vlink)#` prompt.

The OSPF interface parameters for this process are automatically set to their default values for virtual links. You can optionally adjust the following OSPF virtual link settings.

1. Set the time interval between OSPF hello packets for the OSPF virtual links using the following command. For command details, see hello-interval .

   `hello-interval <seconds>`

For example:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink)# hello-interval 30
```

2. Set the interval after which a neighbor is declared dead if no hello packet is received on the OSPF virtual links. Use the following command. For command details, see dead-interval
   `dead-interval <seconds>`
3. Set the time between retransmitting lost link state advertisements for the OSPF virtual links using the following command. For command details, see retransmit-interval.
   `retransmit-interval <seconds>`
4. Sets the transit delay in Link state transmission for the OSPF virtual links using the following command. For command details, see transit-delay.
   `transit-delay <seconds>`

Configuring OSPFv3 interface settings

Prerequisites

You must be in the interface configuration context, as indicated by the `switch(config-if)#` prompt.

You can optionally adjust the following OSPFv3 interface settings.

1. Set the interface cost using the following command. For command details, see ipv6 ospfv3 cost.
   `ipv6 ospfv3 cost <interface-cost>`
   For example, the following command sets the cost associated with interface 1/1/1 to 100.

```
switch(config)# interface 1/1/1
switch(config-if)# ipv6 ospfv3 cost 100
```

2. Set the time interval between OSPFv3 hello packets for the OSPFv3 interface using the following command. For command details, see ipv6 ospfv3 hello-interval.
   `ipv6 ospfv3 hello-interval <seconds>`
3. Set the interval after which a neighbor is declared dead if no hello packet is received on the OSPFv3 interface using the following command. For command details, see ipv6 ospfv3 dead-interval.
   `ipv6 ospfv3 dead-interval <seconds>`
4. Set the time between re-transmitting lost link state advertisements for the OSPFv3 interface using the following command. For command details, see ipv6 ospfv3-af retransmit-interval.
   `ipv6 ospfv3 retransmit-interval <seconds>`
5. Sets the time delay in Link state transmission for the OSPFv3 interface using the following command. For command details, see ipv6 ospfv3 transit-delay.
   `ipv6 ospfv3 transit-delay <seconds>`
6. Set the OSPFv3 network type for the interface. For command details, see ipv6 ospfv3 network.
   `ipv6 ospfv3 network {broadcast|point-to-point}`
7. Set the OSPFv3 priority on the interface using the following command. For command details, see ipv6 ospfv3 priority.
   `ipv6 ospfv3 priority <number-value>`

8. Set the OSPFv3 interface as OSPFv3 passive interface. With this setting the interface participates in OSPFv3 but does not send or receive packets on that interface. For command details, see ipv6 ospfv3 passive.
   ```
   ipv6 ospfv3 passive
   ```

Configuring BFD for OSPFv3

Enabling BFD Support for OSPFv3 over IPv6 enables OSPFv3 to register with Bidirectional Forwarding Detection (BFD) to receive forwarding path detection failure messages. You can either configure BFD for OSPFv3 globally on all interfaces or configure it selectively on one or more interfaces. BFD creates a session in asynchronous mode as soon as the switch reaches the 2-Way state with a neighbor. Once the session is established BFD will commence sending echos for path failure detection.

There are two methods for enabling BFD for OSPFv3:

1. Enable BFD for all interfaces enabled for OSPFv3 by using the BFD all-interfaces command in router configuration mode
2. Enable BFD for a subset of interfaces that have OSPFv3 enabled by using the IPv6 ospfv3 BFD command in interface configuration mode

OSPFv3 needs to be enabled prior to enabling BFD on one or more interfaces.

Examples

Enable BFD on all OSPFv3 Interfaces

```
switch(config)# router ospfv3 1
                switch(config-ospfv3-1)# bfd all-interfaces
```

Disable BFD on all OSPFv3 Interfaces

```
switch(config)# router ospfv3 1
                switch(config-ospfv3-1)# no bfd all-interfaces
```

Enable BFD on a single OSPFv3 interface

```
switch(config)# interface 1/1/1
                switch(config-if)# ipv6 ospfv3 bfd
```

Disable BFD on a single OSPFv3 interface

```
switch(config)# interface 1/1/1
                switch(config-if)# ipv6 ospfv3 bfd disable
```

Set BFD as default on an OSPFv3 interface

```
switch(config)# interface 1/1/1
                switch(config-if)# no ipv6 ospfv3 bfd
```

Configuring all OSPFv3 interfaces as passive

Prerequisites

You must be in the OSPFv3 router configuration context, as indicated by the `switch(config-ospfv3-1)#` prompt.

OSPFv3 sends LSAs to all other routers in the same AS. To limit the flooding of LSAs throughout the AS, you can configure OSPFv3 to be passive.

Procedure

Configure all OSPFv3 interfaces as passive using the following command. For command details, see passive-interface default.

`passive-interface default`

```
switch(config-ospfv3-1)# passive-interface default
```

Configuring SPF throttling timers

SPF calculation is throttled with default timer values (start-time 200ms, hold-time 1000ms, max-wait-time 5000ms). You can throttle SPF calculation by configuring non-default timers to improve performance of a specific network configuration.

Prerequisites

You must be in the router configuration context, as indicated by the `switch(config-router)#` prompt.

Procedure

Reconfigure SPF throttling timers using the following command. For command details, see timers throttle spf.

`timers throttle spf start-time <milliseconds> hold-time <milliseconds> max-wait-time <milliseconds>`

For example, use the following command to set start-time to 500ms, hold-time to 3000ms and max-wait-time to 9000ms:

```
switch(config-ospf-1)# timers throttle spf start-time 500 hold-time
3000 max-wait-time 9000
```

Viewing OSPFv3 information

Prerequisites

Use these show commands from the Manager context, as indicated by the `switch#` prompt.

Procedure

To view OSPFv3 information, use the following commands. For command details and examples, click the links.

- To view general OSPFv3, area, state and configuration information use: show ipv6 ospfv3.
- To view information about OSPFv3 interfaces use: show ipv6 ospfv3 interface.

- To view information about OSPFv3 neighbors use: show ipv6 ospfv3 neighbors .
- To view OSPFv3 routing table information use: show ipv6 ospfv3 routes.
- To view OSPFv3 statistics for the OSPFv3 interfaces use show ipv6 ospfv3 statistics interface.

Clearing OSPFv3 statistics on a switch

Clear OSPFv3 event statics using the following command. For command details, see clear ipv6 ospfv3 statistics

```
clear ipv6 ospfv3 [<PROCESS-ID>] statistics [interface [<INTERFACE-NAME>]] [all-vrfs | vrf <VRF-NAME>]
```

For example, the following command clears OSPFv3 event statistics from interface 1/1/1.

```
switch# clear ipv6 ospfv3 statistics interface 1/1/1
```

OSPFv3 commands

active-backbone

```
active-backbone stub-default-route
no active-backbone stub-default-route
```

Description

This command enables the router to send a default route to stub areas if there is an active loopback link in the backbone area. The configuration is not required if backbone area has neighbors or passive interfaces configured. By default active backbone detection is enabled.

Examples

```
switch(config)# router ospf 1
switch(config-ospf-1)# active-backbone stub-default-route
```

```
switch(config)# no active-backbone stub-default-route
```

Command History

Release	Modification
10.10.1000	Command Introduced

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400	config-ospf-<PROCESS-ID> config-ospfv3-<PROCESS-ID>	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300 9300S 10000		

area

```
area <AREA-ID>
no area <AREA-ID>
```

Description

Creates a normal area with **<AREA-ID>** set if not present. If area is present and is not the normal area, this command changes the area type to normal area.

The **no** form of this command deletes the area with the **<AREA-ID>** specified. The area can be of any type (stub, stub no-summary, and default normal area).

Parameter	Description
<AREA-ID>	Specifies the area ID is one of the following formats. OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to 4294967295.

Examples

Creating a normal area:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 1
switch(config-ospfv3-1)# area 1.1.1.1
```

Deleting an area:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no area 1
switch(config-ospfv3-1)# no area 1.1.1.1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

area authentication ipsec

```
area <AREA-ID> authentication ipsec spi <SPI-INDEX> <AUTH-TYPE> [<KEY-TYPE> <AUTH-KEY>]
no area <AREA-ID> authentication
```

Description

Configures IPsec AH authentication for the specified area. OSPFv3 interfaces which have IPsec configured at the interface context will not use area level IPsec.

The **no** form of this command removes IPsec AH authentication for the specified area.

IPsec is not supported for 6in6 tunnel interfaces.

Parameter	Description
<AREA-ID>	Specifies the area ID is one of the following formats: OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to 4294967295.
`spi` *<SPI-INDEX>*	Specifies the Security Parameters Index (SPI) to use. The SPI is an identification tag carried in the IPsec AH header. It enables the receiving OSPF process to select and use the Security Association (SA) from the SA table. The SPI must be unique on the switch. Range: 256 to 4294967295.
<AUTH-TYPE>	Specifies the authentication type: **md5** or **sha1**.
<KEY-TYPE>	Specifies the key type to use: **plaintext** (unencrypted), **hex-string** (encrypted) or **ciphertext** (encrypted).
<AUTH-KEY>	Specifies the authentication key.

When the authentication key is not provided on the command line, plaintext key prompting occurs upon pressing Enter. The entered key characters are masked with asterisks.

Examples

Setting area 0 to use IPsec authentication with a provided plaintext authentication key:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 0 authentication ipsec spi 256 sha1 plaintext
F82#450
```

Setting area 5 to use IPsec authentication with a prompted plaintext authentication key:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 5 authentication ipsec spi 256 sha1
Enter the IPsec authentication key: ********
Re-Enter the IPsec authentication key: ********
```

Removing IPsec authentication from area 1:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no area 1 authentication
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

area encryption ipsec

```
area <AREA-ID> encryption ipsec spi <SPI-INDEX> <AUTH-TYPE>
    [<KEY-TYPE> <AUTH-KEY> <ENCR-TYPE> [<KEY-TYPE> <ENCR-KEY>]]
no area <AREA-ID> encryption
```

Description

Configures IPsec ESP with the authentication and encryption algorithm types and keys for the specified area. OSPFv3 interfaces with IPsec configured at the interface context will not use area level IPsec ESP configuration.

The **no** form of this command removes IPsec ESP from the specified area.

IPsec is not supported for 6in6 tunnel interfaces.

Parameter	Description
<AREA-ID>	Specifies the area ID is one of the following formats. OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to 4294967295.
`spi` *<SPI-INDEX>*	Specifies the Security Parameters Index (SPI) to use. The SPI is an identification tag carried in the IPsec AH header. It enables the receiving OSPF process to select and use the Security Association (SA) from the SA table. The SPI must be unique on the switch. Range: 256 to 4294967295.
<AUTH-TYPE>	Specifies the authentication type: **md5** or **sha1**.
<KEY-TYPE>	Specifies the key type to use: **plaintext** (unencrypted), **hex-string** (encrypted) or **ciphertext** (encrypted).
<AUTH-KEY>	Specifies the authentication key.
<ENCR-TYPE>	Specifies the encryption type: **des**, **3des**, **aes**, or **null**. **NOTE:** Encryption type **aes** is considered to be AES128, AES192 or AES256 based on key length.
<ENCR-KEY>	Specifies the encryption key.

When the authentication key is not provided on the command line, plaintext authentication key prompting occurs upon pressing Enter, followed by encryption type prompting, and finally plaintext encryption key prompting. The entered key characters are masked with asterisks.

When the authentication key and encryption type are provided on the command line but the encryption key is not provided, plaintext encryption key prompting occurs upon pressing Enter. The entered key characters are masked with asterisks.

Examples

Setting area 0 to use IPSec ESP with provided authentication and encryption keys:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 0 encryption ipsec spi 256 md5 plaintext F824eva
                            des plaintext F82#450b
```

Setting area 5 to use IPSec ESP with prompted authentication and encryption keys:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 5 encryption ipsec spi 256 md5
```

```
Enter the IPsec authentication key: ********
Re-Enter the IPsec authentication key: ********

Enter the IPsec encryption type (3des/aes/des/null)? des

Enter the IPsec encryption key: ********
Re-Enter the IPsec encryption key: ********
```

Setting area 2 to use IPsec ESP with provided authentication password and encryption type but a prompted encryption key:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)#  area 2 encryption ipsec spi 256 md5 plaintext F82# des
Enter the IPsec encryption key: ********
Re-Enter the IPsec encryption key: ********
```

Setting area 0 to use IPSec ESP with provided plaintext authentication key and null encryption:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 0 encryption ipsec spi 256 md5 plaintext axtw null
```

Removing IPSec ESP from area 0:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no area 0 encryption
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

area nssa

```
area <AREA-ID> nssa [no-summary]
no area <AREA-ID> nssa [no-summary]
```

Description

Creates the NSSA area (Not So Stubby Area) with **<AREA-ID>** if not present. If area is present and not NSSA area, this command changes the area type to NSSA area. If **no-summary** is used, area type will be NSSA No-Summary.

The **no** form of this command clears the NSSA area type. That is, the configured area will be changed to default normal area. The **no area <AREA-ID> nssa no-summary** command enables sending inter-area routes into NSSA, but will not unset the area as NSSA.

Parameter	Description
<AREA-ID>	Specifies the area ID is one of the following formats. OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to 4294967295.
`nssa [no-summary]`	Specifies Not So Stubby Area (NSSA) area type. If area is present and not NSSA area, parameter changes the area type to NSSA area. If **no-summary** is specified, area type will be NSSA No-Summary, which means do not inject inter-area routes into NSSA.

Examples

Creating an NSSA area for OSPFv3:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 1 nssa
switch(config-ospfv3-1)# area 1 nssa no-summary
```

Clearing the NSSA area for OSPFv3:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no area 1 nssa
switch(config-ospfv3-1)# no area 1 nssa no-summary
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325P 8360 9300 9300S 10000		

area range

```
area <AREA-ID> range <IP-PREFIX> type {inter-area | nssa} [no-advertise]
no area <AREA-ID> range <IP-PREFIX> type {inter-area | nssa} [no-advertise]
```

Description

Summarizes the routes with the matching address or masks for OSPFv3. This command only works for border routers.

The **no** form of this command unsets the route summarization for the configured IPv4 prefix address on the ABR. When using the **no** form of the command with the **no-advertise** option, enables advertising this range to other areas.

Parameter	Description
<AREA-ID>	Specifies the area ID is one of the following formats. OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to 4294967295.
`range <IP-PREFIX>`	Specifies summarizing routes matching the area range prefix/mask.
`type {inter-area | nssa}`	Specifies the type this address aggregation applies to as either inter-area range prefix or NSSA range prefix.
`no-advertise`	Specifies the address range status as **DoNotAdvertise** (do not advertise this range to other areas).

Examples

Summarizing inter-area or NSSA paths on OSPFv3:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 1 range fd00::/64 type inter-area
switch(config-ospfv3-1)# area 1 range fd00::/64 type nssa
switch(config-ospfv3-1)# area 1 range fd00::/64 type inter-area no-advertise
```

Unsetting summarization on OSPFv3:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no area 1 range fd00::/64 type inter-area
switch(config-ospfv3-1)# no area 1 range fd00::/64 type nssa
switch(config-ospfv3-1)# no area 1 range fd00::/64 type inter-area no-advertise
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

area stub

```
area <AREA-ID> stub [no-summary]
no area <AREA-ID> stub [no-summary]
```

Description

Creates the stub area with **<AREA-ID>** if not present. If area is present and is not the stub area, this command changes the area type to stub area. If **no-summary** is used, area type will be totally stubby area.

The **no** form of this command unsets the area type as stub. The configured area will be changed to the default normal area. The **no area <AREA-ID> stub no-summary** command will start sending Area Border Router (ABR) summary link advertisements into the stub area, but will not unset the stub area.

ABR does not inject the default route in a Totally Stubby Area with loopback in Area 0.0.0.0. As a workaround, configure a passive interface or active neighbors in the backbone area.

Parameter	Description
<AREA-ID>	Specifies the area ID is one of the following formats. OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to 4294967295.
`stub [no-summary]`	Specifies stub area type. If area is present and not stub area, this parameter changes the area type to stub area. If **no-summary** is specified, area type will be totally stubby area, which means do not inject interarea routes into stub.

Examples

Creating a stub area:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 1 stub
switch(config-ospfv3-1)# area 1 stub no-summary
```

Unsetting the stub area type:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1) # no area 1 stub
switch(config-ospfv3-1) # no area 1 sub no-summary
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

area virtual-link

```
area <AREA-ID> virtual-link <ROUTER-ID>
no area <AREA-ID> virtual-link <ROUTER-ID>
```

Description

Creates an OSPF virtual link with remote ABR (if not created already) and enters the vlink context.

The **no** form of this command deletes an OSPF virtual link with the specified router ID of the remote ABR. If no**<ROUTER-ID>** is specified, the **no** form of the command sets the virtual link to the default settings.

Parameter	Description
<AREA-ID>	Specifies the area ID is one of the following formats. OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.

Parameter	Description
	OSPF area identifier in decimal format. Range: 0 to 4294967295.
`virtual-link <ROUTER-ID>`	Configures a virtual link with the specified router ID of the remote ABR.

Examples

Configuring OSPF virtual links:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 100 virtual-link 100.0.1.1
```

Deleting OSPF virtual links:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no area 100 virtual-link 100.0.1.1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

area default-metric

```
area <AREA-ID> default-metric <METRIC>
no area <AREA-ID> default-metric
```

Description

Sets the cost of default-summary LSAs announced to the stub/nssa areas.

The **no** form of this command resets the cost of the default-summary LSAs announced to stub/nssa areas to the default of 1.

Parameter	Description
<AREA-ID>	Specifies the area ID is one of the following formats. OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to 4294967295.
`default-metric` *<METRIC>*	Specifies the default metric of default-summary LSAs announced to the stub/nssa areas, to the specified value. Default: 1. Range: 0 to 16777215.

Examples

Setting cost for default LSA summary:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 1 default-metric 2
switch(config-ospfv3-1)# area 1.1.1.1 default-metric 2
```

Setting cost for default LSA summary to default:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no area 1 default-metric
switch(config-ospfv3-1)# no area 0.0.0.1 default-metric
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-`*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

authentication ipsec

`authentication ipsec spi` *<SPI-INDEX>* *<AUTH-TYPE>* [*<KEY-TYPE>* *<AUTH-KEY>*]
`no authentication`

Description

Configures IPsec AH authentication for the selected Vlink.

The **no** form of this command removes IPsec AH authentication for the selected Vlink.

Parameter	Description
`spi <SPI-INDEX>`	Specifies the Security Parameters Index (SPI) to use. The SPI is an identification tag carried in the IPsec AH header. It enables the receiving OSPF process to select and use the Security Association (SA) from the SA table. The SPI must be unique on the switch. Range: 256 to 4294967295.
<AUTH-TYPE>	Specifies the authentication type: **md5** or **sha1**.
<KEY-TYPE>	Specifies the key type to use: **plaintext** (unencrypted), **hex-string** (encrypted) or **ciphertext** (encrypted).
<AUTH-KEY>	Specifies the authentication key.

When the authentication key is not provided on the command line, plaintext key prompting occurs upon pressing Enter. The entered key characters are masked with asterisks.

Examples

Setting area 1 to use IPsec AH authentication for Vlink with provided plaintext key:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 1
switch(config-ospfv3-1)# area 1 virtual-link 3.3.3.3
switch (config-router-vlink6)# authentication ipsec spi 256 sha1 plaintext F82#450
```

Setting area 1 to use IPsec AH authentication for Vlink with prompted plaintext key:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 1
switch(config-ospfv3-1)# area 1 virtual-link 3.3.3.3
switch (config-router-vlink6)# authentication ipsec spi 256 sha1
Enter the IPsec authentication key: ********
Re-Enter the IPsec authentication key: *******
```

Removing IPsec AH authentication for Vlink on area 1:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 1 virtual-link 3.3.3.3
switch(config-router-vlink6)# no authentication
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-router-vlink`	Administrators or local user group members with execution rights for this command.

clear ipv6 ospfv3 neighbors

`clear ipv6 ospfv3 [<PROCESS-ID>] neighbor [<NEIGHBOR>] [interface [<INTERFACE-NAME>]] [all-vrfs | vrf <VRF-NAME>]`

Description

Resets the neighbor and clears the OSPF neighbor information.

Parameter	Description
<PROCESS-ID>	Specifies the OSPFv3 process ID to clear the statistics for the particular OSPFv3 process. Range: 1 to 65535.
<NEIGHBOR>	Specifies the router ID of a neighbor.
<INTERFACE-NAME>	Specifies the OSPFv3 statistics to clear for the specified interface.
`all-vrfs`	Select to clear the OSPFv3 statistics for all VRFs.
`vrf <VRF-NAME>`	Specifies the name of a VRF.

Example

Clearing the OSPFv3 neighbor information:

```
switch# clear ipv6 ospfv3 1 neighbor
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ipv6 ospfv3 1 neighbor 3.3.3.3
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ipv6 ospfv3 1 neighbor interface 1/1/1
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ipv6 ospfv3 1 neighbor 3.3.3.3 vrf red
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ipv6 ospfv3 neighbor
Performing clear ospf neighbor may result in traffic disruption.
```

```
Do you want to continue (y/n)? y
switch# clear ipv6 ospfv3 neighbor 5.5.5.5
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ipv6 ospfv3 neighbor interface 1/1/1
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ipv6 ospfv3 neighbor 5.5.5.5 vrf red
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

clear ipv6 ospfv3 statistics

```
clear ipv6 ospfv3 [<PROCESS-ID>] statistics [interface [<INTERFACE-NAME>]]
[all-vrfs | vrf <VRF-NAME>]
```

Description

Clears the OSPFv3 event statistics.

Parameter	Description
<PROCESS-ID>	Specifies the OSPFv3 process ID to clear the statistics for the particular OSPFv3 process. Range: 1 to 65535.
<INTERFACE-NAME>	Specifies the OSPFv3 statistics to clear for the specified interface.
`all-vrfs`	Select to clear the OSPFv3 statistics for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF.

Example

Clearing the OSPFv3 event statistics:

```
switch# clear ipv6 ospfv3 statistics
switch# clear ipv6 ospfv3 statistics interface 1/1/1
switch# clear ipv6 ospfv3 statistics interface 1/1/1 vrf vrf_red
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

clear ospfv3-af ipv4 neighbor

```
clear ospfv3-af [<PROCESS-ID>] ipv4 neighbor [<NEIGHBOR-ID>] [interface <INTERFACE-NAME>
| vlan <VLAN-ID> | lag <LAG-ID>] | loopback <LOOPBACK-ID> | tunnel <TUNNEL-ID>] [all-vrfs
| vrf <VRF-NAME>]
```

Description

Resets the neighbor and clears the OSPFv3 AF neighbor information.

Parameter	Description
<PROCESS-ID>	Specifies the OSPFv3 AF process ID to clear the statistics for the particular OSPFv3 AF process. Range: 1 to 65535.
neighbor *<NEIGHBOR-ID>*	Specifies the router ID of a neighbor.
interface *<INTERFACE-NAME>*	Specifies the OSPFv3 AF statistics to clear for the specified interface.
vlan *<VLAN-ID>*	Specifies the vlan interface name.
lag *<LAG-ID>*	Specifies the lag interface name.
loopback *<LOOPBACK-ID>*	Specifies the loopback interface name.
tunnel *<TUNNEL-ID>*	Specifies the tunnel interface name.
all-vrfs	Select to clear the OSPFv3 AF statistics for all VRFs.
vrf *<VRF-NAME>*	Specifies the name of a VRF.

Example

Clearing the OSPFv3 AF neighbor information:

```
switch# clear ospfv3-af 1 ipv4 neighbor
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af 1 ipv4 neighbor 3.3.3.3
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af 1 ipv4 neighbor interface 1/1/1
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af 1 ipv4 neighbor 3.3.3.3 vrf red
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af ipv4 neighbor
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af ipv4 neighbor 5.5.5.5
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af ipv4 neighbor interface 1/1/1
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af ipv4 neighbor 5.5.5.5 vrf red
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

clear ospfv3-af ipv4 statistics

`clear ospfv3-af [<PROCESS-ID>] ipv4 statistics [interface [interface <INTERFACE-NAME> | vlan <VLAN-ID> | lag <LAG-ID>] | loopback <LOOPBACK-ID> | tunnel <TUNNEL-ID>] [all-vrfs | vrf <VRF-NAME>]`

Description

Clears the OSPFv3 AF event statistics.

Parameter	Description
<PROCESS-ID>	Specifies the OSPFv3 AF process ID to clear the statistics for the particular OSPFv3 AF process. Range: 1 to 65535.
`interface` *<INTERFACE-NAME>*	Specifies the OSPFv3 AF statistics to clear for the specified interface.
`vlan` *<VLAN-ID>*	Specifies the vlan interface name.
`lag` *<LAG-ID>*	Specifies the lag interface name.
`loopback` *<LOOPBACK-ID>*	Specifies the loopback interface name.
`tunnel` *<TUNNEL-ID>*	Specifies the tunnel interface name.
`all-vrfs`	Select to clear the OSPFv3 AF statistics for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF.

Example

Clearing the OSPFv3 AF event statistics:

```
switch# clear ospfv3-af ipv4 statistics
switch# clear ospfv3-af ipv4 statistics interface 1/1/1
switch# clear ospfv3-af ipv4 statistics interface 1/1/1 vrf vrf_red
switch# clear ospfv3-af ipv4 statistics interface tunnel1
switch# clear ospfv3-af ipv4 statistics interface tunnel1 vrf vrf_red
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

clear ospfv3-af ipv6 neighbor

`clear ospfv3-af [`*<PROCESS-ID>*`] ipv6 neighbor [`*<NEIGHBOR-ID>*`] [interface` *<INTERFACE-NAME>* `| vlan` *<VLAN-ID>* `| lag` *<LAG-ID>*`] | loopback` *<LOOPBACK-ID>* `| tunnel` *<TUNNEL-ID>*`] [all-vrfs | vrf` *<VRF-NAME>*`]`

Description

Resets the neighbor and clears the OSPFv3 AF neighbor information.

Parameter	Description
<PROCESS-ID>	Specifies the OSPFv3 AF process ID to clear the statistics for the particular OSPFv3 AF process. Range: 1 to 65535.
`neighbor` *<NEIGHBOR-ID>*	Specifies the router ID of a neighbor.
`interface` *<INTERFACE-NAME>*	Specifies the OSPFv3 AF statistics to clear for the specified interface.
`vlan` *<VLAN-ID>*	Specifies the vlan interface name.
`lag` *<LAG-ID>*	Specifies the lag interface name.
`loopback` *<LOOPBACK-ID>*	Specifies the loopback interface name.
`tunnel` *<TUNNEL-ID>*	Specifies the tunnel interface name.
`all-vrfs`	Select to clear the OSPFv3 AF statistics for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF.

Example

Clearing the OSPFv3 AF neighbor information:

```
switch# clear ospfv3-af 1 ipv6 neighbor
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af 1 ipv6 neighbor 3.3.3.3
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af 1 ipv6 neighbor interface 1/1/1
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af 1 ipv6 neighbor 3.3.3.3 vrf red
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af ipv6 neighbor
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af ipv6 neighbor 5.5.5.5
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af ipv6 neighbor interface 1/1/1
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af ipv6 neighbor 5.5.5.5 vrf red
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

clear ospfv3-af ipv6 statistics

```
clear ospfv3-af [<PROCESS-ID>] ipv6 statistics [interface [interface <INTERFACE-NAME> |
vlan <VLAN-ID> | lag <LAG-ID>] | loopback <LOOPBACK-ID> | tunnel <TUNNEL-ID>] [all-vrfs |
vrf <VRF-NAME>]
```

Description

Clears the OSPFv3 AF event statistics.

Parameter	Description
<PROCESS-ID>	Specifies the OSPFv3 AF process ID to clear the statistics for the particular OSPFv3 AF process. Range: 1 to 65535.
`interface` *<INTERFACE-NAME>*	Specifies the OSPFv3 AF statistics to clear for the specified interface.
`vlan` *<VLAN-ID>*	Specifies the vlan interface name.
`lag` *<LAG-ID>*	Specifies the lag interface name.
`loopback` *<LOOPBACK-ID>*	Specifies the loopback interface name.
`tunnel` *<TUNNEL-ID>*	Specifies the tunnel interface name.
`all-vrfs`	Select to clear the OSPFv3 AF statistics for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF.

Example

Clearing the OSPFv3 AF event statistics:

```
switch# clear ospfv3-af ipv6 statistics
switch# clear ospfv3-af ipv6 statistics interface 1/1/1
```

```
switch# clear ospfv3-af ipv6 statistics interface 1/1/1 vrf vrf_red
switch# clear ospfv3-af ipv6 statistics interface tunnel1
switch# clear ospfv3-af ipv6 statistics interface tunnel1 vrf vrf_red
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

dead-interval

```
dead-interval <INTERVAL>
no dead-interval
```

Description

Sets the interval after which a neighbor is declared dead if no hello packet comes in for virtual links.

The **no** form of this command sets the dead interval to default for virtual links. The default value is 40 seconds (generally four times the hello packet interval).

Parameter	Description
<INTERVAL>	Specifies the time interval for the dead interval, in seconds. Range: 1 to 65535. Default: 40.

Examples

Setting OSPFv3 virtual links dead interval:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink6)# dead-interval 30
```

Setting OSPFv3 virtual links dead interval to default:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink6)# no dead-interval
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-router-vlink`	Administrators or local user group members with execution rights for this command.

default-metric

```
default-metric <METRIC-VALUE>
no default-metric
```

Description

Sets the default metric for redistributed routes in the OSPFv3.

The **no** form of this command sets the default metric to be used for redistributed routes into OSPFv3 to the default of 25.

Parameter	Description
<METRIC-VALUE>	Specifies the default metric value to use for redistributed routes. Range: 1 to 1677214. Default: 25.

Examples

Setting default metric for redistributed routes:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# default-metric 36
```

Setting default metric for redistributed routes to the default:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no default-metric
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

disable

```
disable
```

Description

Disables the OSPFv3 process. By default OSPFv3 process is enabled.

This command does not remove the OSPFv3 configurations.

Example

Disabling OSPFv3 process:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# disable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

distance

```
distance [<DISTANCE-VAL> | intra-area [<DISTANCE-VAL>] | inter-area [<DISTANCE-VAL>] |
external [<DISTANCE-VAL>]]
no distance [<DISTANCE-VAL> | intra-area [<DISTANCE-VAL>] | inter-area [<DISTANCE-VAL>] |
external [<DISTANCE-VAL>]]
```

Description

Defines an administrative distance for OSPFv3. Administrative distance is used as a criteria to select the best route when the same route is learned by multiple routing protocols.

The **no** form of this command sets the OSPFv3 administrative distance to the default value of 110. Optionally, administrative distance can be set to default for the specific OSPF route type: intra-area, inter-area, or external type-5 and type-7 routes.

Parameter	Description
<DISTANCE-VAL>	Specifies the OSPFv3 administrative distance. Range: 1 to 255. Default: 110.
`intra-area`	Specifies the OSPFv3 distance for intra-area routes.
`inter-area`	Specifies the OSPFv3 distance for inter-area routes.
`external`	Specifies the OSPFv3 distance for external type 5 and type 7 routes.

Usage

Within a given OSPF process, intra-area routes are always given precedence even when distances are configured for inter-area or external type routes.

Examples

Setting OSPFv3 administrative distance:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# distance 100
switch(config-ospfv3-1)# distance intra-area 24 external 55 inter-area 66
switch(config-ospfv3-1)# distance intra-area 24 external 55
```

```
switch(config-ospfv3-1)# distance external 55
switch(config-ospfv3-1)#exit

switch(config)# router ospfv3 2
switch(config-ospfv3-2)# distance 200
switch(config-ospfv3-2)# distance external 60
switch(config-ospfv3-2)# distance intra-area 24 inter-area 66
```

Setting OSPFv3 administrative distance to the default:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no distance
switch(config-ospfv3-1)# no distance external
switch(config-ospfv3-1)# no distance intra-area
switch(config-ospfv3-1)# no distance inter-area
switch(config-ospfv3-1) # no distance 1
switch(config-ospfv3-1)#exit

switch(config)# router ospfv3 2 vrf blue
switch(config-ospfv3-2)# no distance 200
switch(config-ospfv3-2)# no distance external 60
switch(config-ospfv3-2)# no distance intra-area 24 inter-area 66
```

Command History

Release	Modification
10.14	Added capability to have individual admin distance for multiple OSPF processes in a VRF.
10.09	Added parameters: **intra-area**, **inter-area**, **external**
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

distribute-list prefix

```
distribute-list prefix <prefix-list-name> {in | out}
no distribute-list prefix <prefix-list-name> {in | out}
```

Description

This command uses an existing prefix list to filter routes that are being installed in the routing table or redistributed to another routing protocol.

The **distribute-list prefix** command filters routes in the inbound or the outbound direction. When this command is issued with the **in** parameter, it filters routes from being installed in the routing table, it does not filter LSAs. When this command is issued with the **out** parameter, it filters only the desired redistributed routes from other protocols.

This command requires that your prefix list is already defined using the ipv6 prefix commands. Route-maps are not supported with the distribute-list feature.

Parameter	Description
`prefix <prefix-list-name>`	Specify the name of an existing prefix.
`{in | out}`	Select one of the following parameters to set the filter direction: ▪ **in**: Filter incoming routes into the routing table ▪ **out**: Filter outgoing routing updates

Examples

The following commands enable the filtering of OSPFv3 routes in an IPv6 network, so routes are no longer installed in the routing table or redistributed from another routing protocol.

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# distribute-list prefix listA in
switch(config-ospfv3-1)# distribute-list prefix listB out
```

The following command disables the filtering of OSPFv3 routes in an IPv6 network, so routes can be installed in the routing table or redistributed from another routing protocol.

```
switch# configure terminal
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no distribute-list prefix listA in
switch(config-ospfv3-1)# no distribute-list prefix listB out
```

Command History

Release	Modification
10.13	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
6400 8100 8325 8325H 8325P 8360 9300 9300S 10000		

enable

```
enable
```

Description

Enables OSPFv3 process when disabled. By default OSPFv3 process is enabled.

Example

Enabling OSPFv3 process:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# enable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

encryption ipsec

```
encryption ipsec spi <SPI-INDEX> <AUTH-TYPE> [<KEY-TYPE> <AUTH-KEY>
   <ENCR-TYPE> [<KEY-TYPE> <ENCR-KEY>]]
no encryption
```

Description

Configures IPSec ESP authentication and encryption for the selected Vlink.

The **no** form of this command removes IPSec ESP authentication and encryption for the selected Vlink.

Parameter	Description
`spi <SPI-INDEX>`	Specifies the Security Parameters Index (SPI) to use. The SPI is an identification tag carried in the IPsec AH header. It enables the receiving OSPF process to select and use the Security Association (SA) from the SA table. The SPI must be unique on the switch. Range: 256 to 4294967295.
<AUTH-TYPE>	Specifies the authentication type: **md5** or **sha1**.
<KEY-TYPE>	Specifies the key type to use: **plaintext** (unencrypted), **hex-string** (encrypted) or **ciphertext** (encrypted).
<AUTH-KEY>	Specifies the authentication key.
<ENCR-TYPE>	Specifies the encryption type: **des**, **3des**, **aes**, or **null**. **NOTE:** Encryption type **aes** is considered to be AES128, AES192, or AES256 based on key length.
<ENCR-KEY>	Specifies the encryption key.

When the authentication key is not provided on the command line, plaintext authentication key prompting occurs upon pressing Enter, followed by encryption type prompting, and finally plaintext encryption key prompting. The entered key characters are masked with asterisks.

When the authentication key and encryption type are provided on the command line but the encryption key is not provided, plaintext encryption key prompting occurs upon pressing Enter. The entered key characters are masked with asterisks.

Examples

Setting area 1 to use IPSec ESP authentication and encryption for Vlink with provided plaintext keys:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 1
switch(config-ospfv3-1)# area 1 virtual-link 3.3.3.3
switch(config-router-vlink6)# encryption ipsec spi 256 md5 plaintext F82#
                              des plaintext Plane#88
```

Setting area 1 to use IPSec ESP authentication and encryption for Vlink with prompted plaintext keys and encryption type:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 1
switch(config-ospfv3-1)# area 1 virtual-link 3.3.3.3
switch(config-router-vlink6)# encryption ipsec spi 256 md5
```

```
Enter the IPsec authentication key: ********
Re-Enter the IPsec authentication key: *******

Enter the IPsec encryption type (3des/aes/des/null)? des

Enter the IPsec encryption key: ********
Re-Enter the IPsec encryption key: ********
```

Setting area 1 to use IPSec ESP authentication and encryption for Vlink provided plaintext authentication key and encryption type but prompted plaintext encryption key:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 1
switch(config-ospfv3-1)# area 1 virtual-link 3.3.3.3
switch(config-router-vlink6)# encryption ipsec spi 256 md5 plaintext Fx des
Enter the IPsec encryption key: ********
Re-Enter the IPsec encryption key: ********
```

Setting area 1 to use IPSec ESP authentication for Vlink with a provided plaintext authentication key and null encryption:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 1
switch(config-ospfv3-1)# area 1 virtual-link 3.3.3.3
switch(config-router-vlink6)# encryption ipsec spi 256 md5 plaintext Fx null
```

Removing IPSec ESP authentication and encryption for Vlink on area 1:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 1 virtual-link 3.3.3.3
switch(config-router-vlink6)# no encryption
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300	config-router-vlink	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
9300S 10000		

default-information originate

```
default-information originate [metric <METRIC-VALUE>]
no default-information originate [metric <METRIC-VALUE>]
```

Description

Configures OSPFv3 to advertise the default route (::/0) to its neighbors if it is present in the routing table. Optionally, the metric value can be set for default route ::/0. The default value is 1.

The **no** form of this command disables advertisement of the default route.

Parameter	Description
metric <METRIC-VALUE>	Specifies the OSPF metric value for the default route. Optional. Default: 1.

Examples

Setting advertisement of the default route:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# default-information originate
```

Disabling advertisement of the default route:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no default-information originate
```

Setting advertisement of the default route and specifying an optional metric value of 20:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# default-information originate
switch(config-ospfv3-1)# default-information originate metric 20
```

Disabling advertisement of the default route and setting metric to the default value:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no default-information originate metric
```

Command History

Release	Modification
10.09	Added parameter: **metric <METRIC-VALUE>**
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospf-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

default-information originate always

```
default-information originate always [metric <METRIC-VALUE>]
no default-information originate always [metric <METRIC-VALUE>]
```

Description

Configures OSPFv3 to advertise the default route (::/0) to its neighbors, regardless if it is present in the routing table or not. Optionally, metric can be set for default route ::/0. The default value is 1.

The **no** form of this command disables advertisement of the default route.

Parameter	Description
metric <METRIC-VALUE>	Specifies the OSPFv3 metric value for the default route. Default: 1.

Examples

Setting advertisement of the default route:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# default-information originate always
```

Disabling advertisement of the default route:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no default-information originate always
```

Setting advertisement of the default route with metric set to 20:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# default-information originate always metric 20
```

Disabling advertisement of the default route and setting the metric to the default value:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no default-information originate always metric
```

Command History

Release	Modification
10.09	Added parameter: **metric <METRIC-VALUE>**
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospf-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

graceful-restart

```
graceful-restart
   restart-interval <INTERVAL>
   helper [strict-lsa-check]
   ignore-lost-interface
   no...
```

Description

Configures graceful restart for OSPFv3. By default graceful restart is enabled on the OSPFv3 router.

The **no** form of this command sets the restart interval to the default of 120 seconds or disables helper mode depending on the specified parameters.

Parameter	Description
`restart-interval <INTERVAL>`	Specifies the time another router waits for this router to gracefully restart and selects the maximum time to wait in seconds. Range: 5 to 1800. Default: 120.
`helper`	Specifies that the router will participate in the graceful restart of a neighbor router.
`strict-lsa-check`	(Optional). Use with the **helper** parameter to enable strict Link state Advertisement (LSA) checking when acting as a restart helper for a restarting peer.

Parameter	Description
	NOTE: OSPF neighbors must disable strict LSA checking. If the local node has fewer OSPF interfaces after restarting, then the neighbors that were adjacent on those interfaces will clear up their adjacencies to the restarting node and will send out link state updates to advertise the dropped adjacency. If strict LSA checking is enabled, the restarting router's neighbors will exit helper mode when they receive the updated LSAs and the graceful restart will still fail.
`ignore-lost-interface`	Enable the restarting router to ignore lost OSPF interfaces during a graceful restart process. This setting should be enabled on a high availability system to ensure a graceful restart completes successfully, even if OSPF-enabled links fail due to High Availability events like a switchover or failover. **NOTE:** Enabling this setting means that the hitless restart procedures do not strictly follow those defined in *RFC 3623, Graceful OSPF Restart*.
`no`	Negate any parameter or return the setting to its default..

Examples

Enabling OSPF graceful restart:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# graceful-restart restart-interval 40
switch(config-ospfv3-1)# graceful-restart helper strict-lsa-check
```

Enabling the switch to ignore lost OSPF interfaces during a graceful restart process:

```
switch(config)# router ospfv3 1
switch (config-ospfv3-1)# graceful-restart ignore-lost-interface
```

Setting the restart interval to default, and disabling helper mode:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no graceful-restart restart-interval
switch(config-ospfv3-1)# no graceful-restart helper
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

hello-interval

```
hello-interval <INTERVAL>
no hello-interval
```

For proper operation, the hello interval must be shorter than the dead interval.

Description

Sets the time interval between OSPF hello packets for virtual links.

The **no** form of this command sets the hello interval to the default value of 10 seconds for virtual links.

Parameter	Description
<INTERVAL>	Specifies the time interval for the hello interval, in seconds. Range: 1 to 65535. Default: 10.

Examples

Setting OSPF virtual links hello interval:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink6)# hello-interval 30
```

Setting OSPF virtual links hello interval to default:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink6)# no hello-interval
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-router-vlink`	Administrators or local user group members with execution rights for this command.

ipv6 ospfv3 area

```
ipv6 ospfv3 <PROCESS-ID> area <AREA-ID>
no ipv6 ospfv3 <PROCESS-ID> area <area-id>
```

Description

Runs the OSPFv3 protocol on the interface for the area specified.

To move an interface to a new area, unmap the existing area and then associate a new area with the interface.

The **no** form of this command disables OSPF on the interface and removes the interface from the area. Interfaces which have an IP address configured on the network or in a subset of the network, stop participating in the OSPF protocol

Parameter	Description
<AREA-ID>	Specifies the area ID is one of the following formats. Area ID in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. Area ID as a decimal value. Range: 0-4294967295.
<PROCESS-ID>	Specifies the OSPFv3 process ID. Range: 1 to 65535.

Examples

Setting OSPFv3 network for the area:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ipv6 ospfv3 1 area 1
switch(config-if-vlan)# ipv6 ospfv3 1 area 0.0.0.1
```

Disabling OSPFv3 network for the area:

```
switch(config)# interface 1/1/1
switch(config-if-vlan)# no ipv6 ospfv3 1 area 1
switch(config-if-vlan)# no ipv6 ospfv3 1 area 0.0.0.1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

ipv6 ospfv3 authentication null

```
ipv6 ospfv3 authentication null
```

Description

Configures null authentication on an interface which disables IPsec authentication.

Examples

Disabling IPsec on interface VLAN **1**:

```
switch(config)# interface van 1
switch(config-if-vlan)# ipv6 ospfv3 authentication null
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325P 8360 9300 9300S 10000		

ipv6 ospfv3 authentication ipsec

```
ipv6 ospfv3 authentication ipsec spi <SPI-INDEX> <AUTH-TYPE> [<KEY-TYPE> <AUTH-KEY>]
no ipv6 ospfv3 authentication
```

Description

Configures IPSec AH authentication. OSPFv3 interfaces that have IPsec configured at the interface context will not use area level IPsec.

The **no** form of this command removes IPsec AH authentication for the specified area.

Parameter	Description
`spi <SPI-INDEX>`	Specifies the Security Parameters Index (SPI) to use. The SPI is an identification tag carried in the IPsec AH header. It enables the receiving OSPF process to select and use the Security Association (SA) from the SA table. The SPI must be unique on the switch. Range: 256 to 4294967295.
<AUTH-TYPE>	Specifies the authentication type: **md5** or **sha1**.
<KEY-TYPE>	Specifies the key type to use: **plaintext** (unencrypted), **hex-string** (encrypted) or **ciphertext** (encrypted).
<AUTH-KEY>	Specifies the authentication key.

When the authentication key is not provided on the command line, plaintext key prompting occurs upon pressing Enter. The entered key characters are masked with asterisks.

Examples

Setting interface VLAN **1** to use IPsec authentication with a provided plaintext authentication key:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ipv6 ospfv3 authentication ipsec spi 256 md5 plaintext
F82#
```

Setting interface VLAN **4** to use IPsec authentication with a prompted plaintext authentication key:

```
switch(config)# interface vlan 4
switch(config-if-vlan)# ipv6 ospfv3 authentication ipsec spi 256 md5
Enter the IPsec authentication key: ********
Re-Enter the IPsec authentication key: ********
```

Removing IPsec authentication from interface VLAN **1**:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ipv6 ospfv3 authentication
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

ipv6 ospfv3 cost

```
ipv6 ospfv3 cost <INTERFACE-COST>
no ipv6 ospfv3 cost
```

Description

Sets the cost (metric) associated with a particular interface. The interface cost is used as a parameter to calculate the best routes.

The **no** form of this command sets the cost (metric) associated with a particular interface to the default of 1.

Parameter	Description
<INTERFACE-COST>	Specifies the interface cost value. Range: 1 to 65535. Default: 1.

Examples

Setting OSPFv3 interface cost:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ipv6 ospfv3 cost 100
```

Setting the OSPFv3 interface cost to default:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ipv6 ospfv3 cost
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

ipv6 ospfv3 dead-interval

```
ipv6 ospfv3 dead-interval <INTERVAL>
no ipv6 ospfv3 dead-interval
```

Description

Sets the interval after a neighbor is declared dead when no hello packet is received on the OSPFv3 interface.

The **no** form of this command sets the interval after which a neighbor is declared dead, to the default for the OSPFv3 interface. The default value is 40 seconds (generally four times the hello packet interval).

Parameter	Description
<INTERVAL>	Specifies the time interval for the dead interval, in seconds. Range: 1 to 65535. Default: 40.

Examples

Setting OSPFv3 dead interval on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ipv6 ospfv3 dead-interval 30
```

Setting OSPFv3 dead interval to default on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ipv6 ospfv3 dead-interval
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if-vlan`	Administrators or local user group members with execution rights for this command.

ipv6 ospfv3 encryption ipsec

```
ipv6 ospfv3 encryption ipsec spi <SPI-INDEX> <AUTH-TYPE>
    [<KEY-TYPE> <AUTH-KEY> <ENCR-TYPE> [<KEY-TYPE> <ENCR-KEY>]]
no ipv6 ospfv3 encryption
```

Description

Configures IPsec ESP authentication. OSPFv3 interfaces that have IPsec configured at the interface context will not use area level IPsec ESP.

The **no** form of this command removes IPsec ESP for the specified area.

Parameter	Description
`spi <SPI-INDEX>`	Specifies the Security Parameters Index (SPI) to use. The SPI is an identification tag carried in the IPsec AH header. It enables the receiving OSPF process to select and use the Security Association (SA) from the SA table. The SPI must be unique on the switch. Range: 256 to 4294967295.
<AUTH-TYPE>	Specifies the authentication type: **md5** or **sha1**.
<KEY-TYPE>	Specifies the key type to use: **plaintext** (unencrypted), **hex-string** (encrypted) or **ciphertext** (encrypted).
<AUTH-KEY>	Specifies the authentication key.
<ENCR-TYPE>	Specifies the encryption type: **des**, **3des**, **aes**, or **null**.

Parameter	Description
	NOTE: Encryption type **aes** is considered to be AES128, AES192, or AES256 based on key length.
<ENCR-KEY>	Specifies the encryption key.

When the authentication key is not provided on the command line, plaintext authentication key prompting occurs upon pressing Enter, followed by encryption type prompting, and finally plaintext encryption key prompting. The entered key characters are masked with asterisks.

When the authentication key and encryption type are provided on the command line but the encryption key is not provided, plaintext encryption key prompting occurs upon pressing Enter. The entered key characters are masked with asterisks.

Examples

Setting interface VLAN **1** to use IPsec ESP with provided authentication and encryption keys:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ipv6 ospfv3 encryption ipsec spi 256 sha1 plaintext F82
                        des plaintext F82#450b
```

Setting interface VLAN **3** to use IPsec ESP with prompted authentication and encryption keys:

```
switch(config)# interface vlan 3
switch(config-if-vlan)# ipv6 ospfv3 encryption ipsec spi 256 sha1
Enter the IPsec authentication key: ********
Re-Enter the IPsec authentication key: ********

Enter the IPsec encryption type (3des/aes/des/null)? des

Enter the IPsec encryption key: ********
Re-Enter the IPsec encryption key: ********
```

Setting interface VLAN **4** to use IPsec ESP with provided authentication password and encryption type but a prompted encryption key:

```
switch(config)# interface vlan 4
switch(config-if-vlan)# ipv6 ospfv3 encryption ipsec spi 256 sha1 plaintext F82
des
Enter the IPsec encryption key: ********
Re-Enter the IPsec encryption key: ********
```

Setting interface VLAN **1** to use IPSec ESP with provided plaintext authentication key and null encryption:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ipv6 ospfv3 encryption ipsec spi 256 sha1 plaintext F82
null
```

Removing IPsec from interface VLAN **1**:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ipv6 ospfv3 encryption
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

ipv6 ospfv3 encryption null

`ipv6 ospfv3 encryption null`

Description

Configures NULL ESP on an interface which disables IPsec ESP.

Examples

Disable IPsec ESP on interface VLAN *1*:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ipv6 ospfv3 encryption null
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

ipv6 ospfv3 hello-interval

```
ipv6 ospfv3 hello-interval <INTERVAL>
no ipv6 ospfv3 hello-interval
```

Description

Sets the time interval between OSPFv3 hello packets for the OSPFv3 interface.

The **no** form of this command sets the time interval between OSPFv3 hello packets to the default for the OSPFv3 interface of 10 seconds.

Parameter	Description
<INTERVAL>	Specifies the time interval between hello packets, in seconds. Range: 1 to 65535. Default: 10.

Examples

Setting OSPFv3 hello interval on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ipv6 ospfv3 hello-interval 30
```

Setting OSPFv3 hello interval to default on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ipv6 ospfv3 hello-interval
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

ipv6 ospfv3 network

```
ipv6 ospfv3 network {broadcast|point-to-point}
no ipv6 ospfv3 network
```

Description

Configures the network type for the interface. By default the network type is broadcast network.

The **no** form of this command sets the network type for the interface to the system default of broadcast network.

Parameter	Description
`broadcast`	Specifies the OSPFv3 network type as a broadcast multiaccess network.
`point-to-point`	Specifies the OSPFv3 network type as a point-to-point network.

Examples

Setting OSPFv3 network type for the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ipv6 ospfv3 network broadcast
switch(config-if-vlan)# ipv6 ospfv3 network point-to-point
```

Disabling OSPFv3 network type for the interface to system default of broadcast network:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ipv6 ospfv3 network
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

ipv6 ospfv3 passive

```
ipv6 ospfv3 passive
no ipv6 ospfv3 passive
```

Description

Configures the interface as an OSPFv3 passive interface. With this setting, the interface participates in the OSPF, but does not send or receive OSPF packets on that interface.

The **no** form of this command resets the interface as active. With this setting, the interface starts sending and receiving OSPF packets.

Examples

Setting the interface as OSPFv3 passive interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ipv6 ospfv3 passive
```

Setting the interface as OSPFv3 active interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ipv6 ospfv3 passive
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
6400 8100 8325 8325H 8325P 8360 9300 9300S 10000		

ipv6 ospfv3 priority

```
ipv6 ospfv3 priority <number-value>
no ipv6 ospfv3 priority
```

Description

Sets the OSPFv3 priority for the interface. The larger the numeric value of the priority, the higher the chance it will become the designated router. Setting a priority of 0 makes the router ineligible to become a designated router or back up designated router.

The **no** form of this command sets the OSPFv3 priority for the interface to the default of 1.

Parameter	Description
<number-value>	Specifies the OSPFv3 priority value. Default: 1. Range: 0 to 255.

Examples

Setting the OSPFv3 priority for the interface:

```
switch(config)# interface vlan /1
switch(config-if-vlan)# ipv6 ospfv3 priority 50
```

Setting the OSPFv3 priority for the interface to the default of 1:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ipv6 ospfv3 priority
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A,	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000		

ipv6 ospfv3-af retransmit-interval

```
ipv6 ospfv3-af retransmit-interval <seconds>
no ipv6 ospfv3-af retransmit-interval
```

Description

Sets the time between retransmitting lost link state advertisements for the OSPFv3 interface.

The **no** form of this command sets the time between retransmitting lost link state advertisements to the default 5 seconds.

Parameter	Description
<seconds>	Specifies the time interval for the retransmit interval, in seconds. Range: 1 to 3600. Default: 5

Examples

Setting OSPFv3 retransmit interval on the interface:

```
switch# configure terminal
switch(config)# interface 1
switch(config-if)# ospfv3-af ipv6 retransmit-interval 30
```

Setting OSPFv3 retransmit interval to the default on the interface:

```
switch# configure terminal
switch(config)# interface 1
switch(config-if)# no ospfv3-af ipv6 retransmit-interval
switch(config-if)# no ospfv3-af ipv6 retransmit-interval 30
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

ipv6 ospfv3 shutdown

```
ipv6 ospfv3 shutdown
no ipv6 ospfv3 shutdown
```

Description

Disables OSPFv3 on the interface. The interface state changes to Down. It does not remove the interface from the OSPF area. To remove the interface, use the command **no ip ospf area**.

The **no** form of this command re-enables OSPFv3 on the interface.

Examples

Disabling OSPFv3 on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ipv6 ospfv3 shutdown
```

Re-enabling OSPFv3 on the interface:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ipv6 ospfv3 shutdown
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400	`config-if` `config-if-vlan`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300 9300S 10000		

ipv6 ospfv3 transit-delay

```
ipv6 ospfv3 transit-delay <DELAY>
no ipv6 ospfv3 transit-delay
```

Description

Sets the time delay in Link state transmission for the OSPFv3 interface.

The **no** form of this command sets the transit delay in Link state transmission to the default of 1 second.

Parameter	Description
<DELAY>	Specifies the time delay for the transit delay, in seconds. Range: 1 to 3600. Default: 1.

Examples

Setting OSPFv3 transit delay on the interface

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ipv6 ospfv3 transit-delay 30
```

Setting OSPFv3 transit delay to default on the interface

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ipv6 ospfv3 transit-delay
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400	config-if config-if-vlan	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300 9300S 10000		

maximum-paths

```
maximum-paths <MAXIMUM>
no maximum-paths
```

Description

Sets the maximum number of ECMP routes that OSPFv3 can support.

The **no** form of this command sets the maximum number of ECMP routes that OSPFv3 can support to the default value of 4.

Parameter	Description
<MAXIMUM>	Specifies the maximum number of ECMP routes. Range: 1 to 32. Default: 4.

Examples

Setting maximum number of parallel routes:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# maximum-paths 32
```

Setting maximum number of parallel paths to the default value of 4:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no maximum-paths
```

Command History

Release	Modification
10.10	Increased upper limit of range of **<MAXIMUM>** parameter to 32.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except	config-ospfv3-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000		

max-metric router-lsa

```
max-metric router-lsa [on-startup <INTERVAL>]
no max-metric router-lsa [on-startup]
```

Description

Sets the protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their shortest path first (SPF) calculations. If the on-startup parameter is used, the router is configured to advertise a maximum metric at startup. That is, for the time specified in seconds, or the default value of 600 seconds.

To disable advertisement of the maximum metric, use the **no** form of the command.

The **no** form of this command advertises the normal cost metrics instead of advertising the maximized cost metric. This setting causes the router to be considered in traffic forwarding.

Parameter	Description
`on-startup <INTERVAL>`	Automatically advertises the stub Router-LSA (or maximize the router-LSA cost metric) for a specified time interval upon OSPFv3 startup. Specifies the time in seconds. Range: 5 to 86400. Default: 600.

Examples

Setting to maximize the cost metrics for Router LSA:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# max-metric router-lsa
switch(config-ospfv3-1)# max-metric router-lsa on-startup 3000
```

Setting to advertise the normal cost metrics instead of advertising the maximized cost metric:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no max-metric router-lsa
switch(config-ospfv3-1)# no max-metric router-lsa on-startup
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

passive-interface default

```
passive-interface default
no passive-interface default
```

Description

Sets all OSPFv3 interfaces as passive.

The **no** form of this command sets all the OSPFv3 interfaces as active.

Examples

Setting OSPFv3-enabled interfaces as passive:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# passive-interface default
```

Setting OSPFv3-enabled interfaces as active:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no passive-interface default
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

redistribute

```
redistribute {bgp | connected | host-routes | local loopback | static | ripng | ospf
<PROCESS-ID>}
             [route-map <ROUTE-MAP-NAME>]
no redistribute {bgp | connected | host-routes | local loopback | static | ripng | ospf
<PROCESS-ID>}
                [route-map <ROUTE-MAP-NAME>]
```

Description

Redistributes routes originating from other protocols, or from another OSPFv3 process, to the current OSPFv3 process.

If a route map is specified, then only the routes that pass the match clause specified in the route map are redistributed to OSPFv3. Configuration is not allowed if the referenced route map has not yet been configured.

If you try to redistribute routes from an OSPFv3 process which is not created, you are prompted to allow the OSPFv3 process to be auto-created before proceeding with redistribution. If you confirm at the prompt, the OSPFv3 process is created with defaults and redistribution configuration applied. If you deny at the prompt, redistribution configuration is skipped.

If command **route-redistribute active-routes-only** has been issued, only the routes from other protocols which are selected for forwarding are considered for redistribution into OSPFv3.

The **no** form of this command disables redistribution of routes to the current OSPFv3 process.

Parameter	Description
`bgp`	Specifies redistributing BGP (Border Gateway Protocol) routes.
`connected`	Specifies redistributing connected (directly attached subnet or host).
`host-routes`	Redistributes ARP entries. ARP entries learnt through EVPN or over VXLAN tunnels are not redistributed. The **host-routes** parameter is applicable to HPE Aruba Networking 8325, 10000, 8360, and 9300/9300S Switch series only.
`local loopback`	Specifies redistributing local routes of the loopback interface.
`static`	Specifies redistributing static routes.

Parameter	Description
`ripng`	Specifies redistributing RIPng routes.
`ospf <PROCESS-ID>`	Specifies redistributing routes from the specified OSPFv3 process ID. Range: 1 to 65535.
`route-map <ROUTE-MAP-NAME>`	Specifies redistribution filtering by route map. To create a route map, use command **route-map**.

Examples

Redistributing routes to OSPFv3:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# redistribute bgp
switch(config-ospfv3-1)# redistribute bgp route-map BGP_routes
```

```
switch(config-ospfv3-1)# redistribute host-routes
switch(config-ospfv3-1)# redistribute connected
switch(config-ospfv3-1)# redistribute connected route-map connected_routes
switch(config-ospfv3-1)# redistribute local loopback
switch(config-ospfv3-1)# redistribute local loopback route-map local_routes
switch(config-ospfv3-1)# redistribute static
switch(config-ospfv3-1)# redistribute static route-map static_networks
switch(config-ospfv3-1)# redistribute ripng
switch(config-ospfv3-1)# redistribute ripng route-map rip-routes
switch(config-ospfv3-1)# redistribute ospf 2
```

Redistributing ARP entries:

```
switch(config-ospfv3-1)# redistribute host-routes route-map host_routes12
```

Disabling redistributing routes to OSPFv3:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no redistribute bgp
switch(config-ospfv3-1)# no redistribute bgp route-map BGP_routes
```

```
 switch(config-ospfv3-1)# no redistribute host-routes
switch(config-ospfv3-1)# no redistribute connected
switch(config-ospfv3-1)# no redistribute connected route-map connected_routes
switch(config-ospfv3-1)# no redistribute local loopback
switch(config-ospfv3-1)# no redistribute local loopback route-map local_routes
switch(config-ospfv3-1)# no redistribute static
switch(config-ospfv3-1)# no redistribute static route-map static_networks
switch(config-ospfv3-1)# no redistribute ripng
switch(config-ospfv3-1)# no redistribute ripng route-map rip-routes
switch(config-ospfv3-1)# no redistribute ospf 2
```

Command History

Release	Modification
10.13	Added **route-map** support for **host-routes**.
10.12	Added information about the **host-routes** parameter.
10.08	Added **route-map** support for supported redistribute source-protocols. Updated information and examples.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

reference-bandwidth

```
reference-bandwidth <BANDWIDTH>
no reference-bandwidth
```

Description

Sets the reference bandwidth for OSPFv3. If the OSPFv3 interface cost is not explicitly set, then the cost of all the OSPFv3 interfaces is recalculated based on the reference bandwidth and link speed of the interface.

For VLAN interfaces the calculated link speed value is 1 Gbps (if the OSPFv3 interface cost is not explicitly set).

The **no** form of this command sets the reference bandwidth for OSPF to the default of 100000 Mbps.

Parameter	Description
<BANDWIDTH>	Specifies the reference bandwidth used to calculate the cost of an interface in Mbps. Range: 1 to 4000000. Default: 100000.

Examples

Setting the reference bandwidth:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# reference-bandwidth 40000
```

Setting the reference bandwidth to the default value:

```
switch(config)# routerv3 ospf 1
switch(config-ospfv3-1)# no reference-bandwidth
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

retransmit-interval

```
retransmit-interval <INTERVAL>
no retransmit-interval
```

Description

Sets the time between retransmitting lost link state advertisements for virtual links.

The **no** form of this command sets the time between retransmitting lost link state advertisements to the default of 5 seconds for virtual links.

Parameter	Description
<INTERVAL>	Specifies the time interval for the retransmit interval, in seconds. Range: 1 to 3600. Default: 5.

Examples

Setting OSPFv3 virtual links retransmit interval:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink6)# retransmit-interval 30
```

Setting OSPFv3 virtual links retransmit interval to default:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink6)# no retransmit-interval
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-router-vlink`	Administrators or local user group members with execution rights for this command.

router-id

```
router-id <ROUTER-ADDRESS>
no router-id
```

Description

Sets an ID for the router in an IPv4 address format.

The **no** form of this command unconfigures the router-id for the instance and sets the router-id to the default. The router-id is changed to the dynamically selected router-id. The default router-id 0.0.0.0 updates to the routing stack that triggers to auto-elect a router-id based on the highest IP address of loopback interface, or the highest IP address of interfaces.

Parameter	Description
<ROUTER-ADDRESS>	Specifies the router address in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.

Examples

Setting router-id in the OSPFv3 context:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1) # router-id 1.1.1.1
```

Unconfiguring router-id:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no router-id
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

router ospfv3

```
router ospfv3 <PROCESS-ID> [vrf <VRF-NAME>]
no router ospfv3 <PROCESS-ID> [vrf <VRF-NAME>]
```

Description

Creates the OSPFv3 process (if not created already) and enters the router OSPFv3 instance context. Optionally if specified, you can specify a named VRF, or the default VRF if the <*vrf-name*> is not specified. Only one OSPFv3 process is allowed per VRF.

The **no** form of this command removes the OSPFv3 instance. If a VRF is specified, it removes the OSPF instance from the named VRF, or the default VRF if the <*var-name*> is not specified.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 process ID. Length: 1 to 65535.
vrf *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Examples

Entering the router OSPFv3 instance:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)#
```

Setting the router OSPFv3 VRF instance:

```
switch(config)# router ospfv3 1 vrf vrf_red
```

Removing the router OSPFv3 instance:

```
switch#(config)# no router ospfv3 1
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

show ipv6 ospfv3

`show ipv6 ospfv3 [<PROCESS-ID>] [all-vrfs | vrf <VRF-NAME>]`

Description

Shows OSPFv3 information including area, state, and configuration information.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 process ID optionally to show OSPFv3 information for a particular OSPFv3 process. Range: 1 to 65535.
`all-vrfs`	Select to show OSPFv3 information for all VRFs.
`vrf <VRF-NAME>`	Specifies the name of a VRF. Default: default.

Example

Showing general OSPFv3 configurations:

```
switch# show ipv6 ospfv3 200
VRF : default                              Process : 200
-------------------------------------------------------------------

Router ID              : 1.1.1.1           OSPFv3                 : Enabled
BFD                    : Disabled          SPF Start Interval     : 200 ms
SPF Hold Interval      : 1000 ms           SPF Max Wait Interval  : 5000 ms
LSA Start Interval     : 5000 ms           LSA Hold Interval      : 1000 ms
LSA Max Wait Interval  : 1000 ms           LSA Arrival Interval   : 1000 ms
External LSAs          : 0                 Checksum Sum           : 0
ECMP                   : 4                 Reference Bandwidth    : 100000 Mbps
Area Border            : Yes               AS Border              : No
GR Status              : Enabled           GR Interval            : 120 sec
GR State               : Inactive          GR Exit Status         : None
GR Helper              : Enabled           GR Strict LSA Check    : Enabled
GR Ignore Lost I/F     : Disabled          Internal Process ID    : 1
Summary address:
  prefix fd00::1/64, advertise, tag 10

Area      Total Active
----------------------
Normal      2      2
Stub        0      0

Area  : 0.0.0.0
---------------------------------
Area Type             : Normal       Status               : Active
Total Interfaces      : 1            Active Interfaces    : 1
Passive Interfaces    : 0            Loopback Interfaces  : 0
SPF calculation count : 4
Area ranges:
  fd00::1/64, inter-area, no-advertise
AH Authentication     : SHA1, SPI 256
Number of LSAs        : 5            Checksum Sum         : 99122

Area  : 0.0.0.1
---------------------------------
Area Type             : Normal       Status               : Active
Total Interfaces      : 1            Active Interfaces    : 1
Passive Interfaces    : 0            Loopback Interfaces  : 0

SPF Calculation Count : 4
Area ranges:
  fd00::1/64, inter-area, no-advertise
ESP Authentication    : SHA1
Encryption            : 3DES, SPI 256
Number of LSAs        : 5            Checksum Sum         : 99122
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ipv6 ospfv3 border-routers

```
show ipv6 ospfv3 [<PROCESS-ID>] border-routers [all-vrfs | vrf <VRF-NAME>]
```

Description

Shows the OSPFv3 routing table entries for Area Border Router (ABR) and Autonomous System Border Router (ASBR).

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 process ID to show the OSPFv3 routing table entries for ABR and ASBR for the particular OSPFv3 process. Range: 1 to 65535.
`all-vrfs`	Select to show OSPFv3 border router information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Example

On the HPE Aruba Networking 6400 Switch Series, interface identification differs.

Showing OSPFv3 border routers information for VRF **vrf_red**:

```
switch# show ipv6 ospfv3 border-routers vrf vrf_red
VRF : vrf_red       Process ID : 1
             Internal Routing Table
----------------------------------------------------------

Codes: i - Intra-area route, I - Inter-area route
   Router-ID Cost  Type    Area    SPF  Nexthop                       Interface
i  1.1.1.1   1     ASBR    0.0.0.0 9    fe80::7272:cfff:fe9a:a15d   1/1/2
i  3.3.3.3   1     ASBR    1.1.1.1 9    fe80::7272:cfff:fe1f:d80   ** tunnel1
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ipv6 ospfv3 interface

```
show ipv6 ospfv3 [<PROCESS-ID>] interface [<INTERFACE-NAME>] [brief]
[all-vrfs | vrf <VRF-NAME>]
```

Description

Shows information about OSPFv3 enabled interfaces.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 process ID optionally to show the OSPFv3 enabled interfaces for the particular OSPFv3 process. Range: 1 to 65535.
<INTERFACE-NAME>	Selects to show information only for the specified OSPFv3-enabled interface.
`brief`	Include this parameter to display a brief overview of the following OSPF configuration information. ■ Interface: OSPF interface name. ■ Area: OSPF area ID. ■ Cost: The metric OSPF uses to judge a path's feasibility, calculated as (reference bandwidth / interface bandwidth). ■ State: Indicates if the interface is a designated router (**Dr**) or a backup designated router (**Backup-dr**). ■ Status: Indicates if the interface is **up** or **down**. ■ Flags: P - Passive A - Active.
`all-vrfs`	Select to show OSPF-enabled interface information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Examples

On the HPE Aruba Networking 6400 Switch Series, interface identification differs.

Showing OSPFv3 information for all interfaces in default VRF:

```
switch# show ipv6 ospfv3 interface
Codes: DR - Designated router  BDR - Backup Designated router

Interface 1/3/10 is Up, Line Protocol is Up
```

```
-------------------------------------------
VRF                 : default                          Process             : 1
IPv6 address        : fe80::9020:c203:280a:e800        Area                :
0.0.0.0
Status              : Up                               Network Type        :
Broadcast
Hello Interval      : 10    sec                        Dead Interval       : 40
 sec
Transit Delay       : 1     sec                        Retransmit Interval : 5
 sec
BFD                 : Disabled                         Link Speed          : 1000
Mbps
Cost Configured     : NA                               Cost Calculated     : 100
State/Type          : DR                               Router Priority     : 2
DR                  : 1.1.1.1                          BDR                 :
2.2.2.2
Link LSAs           : 2                                Checksum Sum        :
39245
Authentication      : no                               Passive             : No

Codes: DR - Designated router  BDR - Backup Designated router

Interface 1/3/11 is Up, Line Protocol is Up
-------------------------------------------
VRF                 : default                          Process             : 1
IPv6 address        : fe80::9020:c203:2c0a:e800        Area                :
0.0.0.1
Status              : Up                               Network Type        :
Broadcast
Hello Interval      : 10    sec                        Dead Interval       : 40
 sec
Transit Delay       : 1     sec                        Retransmit Interval : 5
 sec
BFD                 : Disabled                         Link Speed          : 1000
Mbps
Cost Configured     : NA                               Cost Calculated     : 100
State/Type          : BDR                              Router Priority     : 1
DR                  : 3.3.3.3                          BDR                 :
1.1.1.1
Link LSAs           : 2                                Checksum Sum        :
83119
Authentication      : no                               Passive             : No
```

Showing overview information for OSPFv3 enabled interfaces for all VRFs in brief:

```
switch# show ipv6 ospfv3 interface brief all-vrfs
VRF : default                          Process : 1
==================================================

Total Number of Interfaces: 2

Flags: P - Passive  A - Active


Interface     Area               Cost      State              Status      Flags
-------------------------------------------------------------------------------
1/3/10        0.0.0.0            100       DR                 Up          A
1/3/11        0.0.0.1            100       BDR                Up          A
```

Showing overview information for OSPFv3 enabled interfaces for all VRFs:

```
6200(config)# show ipv6 ospfv3 interface br all-vrfs
VRF : default                              Process : 1
================================================

Total Number of Interfaces: 1

Flags: P - Passive  A - Active


Interface     Area               Cost      State              Status      Flags
-------------------------------------------------------------------------------
1/1/1         0.0.0.0            100       BDR                Up          A
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.09	Output of the **show ipv6 ospfv3 interface** command includes flags to indicate whether the interface is in passive or active mode.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ipv6 ospfv3 lsdb

```
show ipv6 ospfv3 [<process-id>] lsdb
   adv-router {<ROUTER-ID>|self}
   area <AREA-ID>
   lsid <link-state-id>
   all-vrfs|vrf <VRF-NAME>}
      as-external
      asbr-summary
      database-summary
      inter-area-prefix
      inter-area-router
      intra-area-prefix
      link
      network
      nssa-external
      router
```

```
    summary
  vsx-peer
```

Description

Shows the OSPFv3 link state database summary for different OSPF LSAs (Link State Advertisement). Use the parameters to get information for a particular LSA.

Parameter	Description
<PROCESS-ID>	Enter an OSPFv3 process ID to display general OSPF information for a particular OSPF process. Range: 1 to 65535.
adv-router {*<ROUTER-ID>*|self}	Select to display link states for a particular advertising router. Specify either a Router ID of the advertising router or specify **self** to show self-originated link states.
area *<AREA-ID>*	Select to display information filtered for the specified area in one of the following formats. OSPFv3 area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPFv3 area identifier in decimal format. Value: 0 to 4294967295.
lsid *<LINK-STATE-ID>*	Select to display information filtered by link state identifier specified in IPv4 address format (A.B.C.D).
all-vrfs|{vrf <vrf-name>	Select **all-vrfs** to display general OSPF information for all VRFs, or use the **vrf *<VRF-NAME>*** option to display information for a specific VRF. Optionally select one of the following parameters to filter the link state database information.
as-external	Show external link states (LSA type 5)
database-summary	Select to display the count of each type of LSA and each area in the database. **NOTE:** The **database-summary** parameter does not support the **area <area-id>**, **lsid <link-state-id>** or **adv-router {<router-id>|self}** parameters.
inter-area-prefix	Show inter-area prefix link states (LSA type 3)
inter-area-router	Show inter-area router link states (LSA type 4)
intra-area-prefix	Show intra-area prefix link states (LSA type 9
link	Show link states (LSA type 8)
network	Show network LSAs (LSA type 2).
nssa-external	Show NSSA external link states (LSA type 7).
router	Show router LSAs (LSA type 1).
vsx-peer	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the

Parameter	Description
	VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Examples

Showing OSPFv3 link state database (LSDB) general information:

```
switch# show ipv6 ospfv3 lsdb
OSPF Router with ID (50.50.50.50) (Process ID 1 VRF default)
============================================================

Router Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID         Link Count  Bits
-----------------------------------------------------------------------------
40.40.40.40     930     0x80000004  0x2ea1      0            3           None
50.50.50.50     935     0x80000002  0x8b52      0            1           E
60.60.60.60     943     0x800003c5  0x9854      0            2           None

Network Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Router Count
-------------------------------------------------------------------
60.60.60.60     944     0x80000001  0x7179      1360007168  2
50.50.50.50     935     0x80000001  0x516a      19          1

Inter Area Prefix Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
-----------------------------------------------------------
40.40.40.40     929     0x80000001  0x2498      131072      FEC0:3344::/32
50.50.50.50     928     0x80000001  0x5b2f      65536       111::/64

Inter Area Router Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Destination Router ID
-------------------------------------------------------------------------
40.40.40.40     929     0x80000001  0x2498      1           33.33.33.33

AS External Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
-----------------------------------------------------------
40.40.40.40     264     0x80000001  0x24cc4     1           10::/64
40.40.40.40     675     0x80000001  0x5b00f     2           11::/64

NSSA External Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
-----------------------------------------------------------
3.3.3.3         264     0x80000001  0x24ac2     1           200::/64

Link-local Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Interface
-------------------------------------------------------------------
50.50.50.50     264     0x80000001  0x653c4     19          1/1/1

Intra Area Prefix Link State Advertisements (Area 0.0.0.0)
```

```
--------------------------------------------------------------------------------
-------------
ADV Router      Age     Seq#        Checksum    LSID        Referenced LS Type
Referenced LSID
--------------------------------------------------------------------------------
-------------
50.50.50.50     263     0x80000001  0x1da34     1           0x2001              0
50.50.50.50     264     0x80000001  0x2a45d     1           0x2002              19
```

Showing AS external link states:

```
switch# show ipv6 ospfv3 lsdb as-external
OSPF Router with ID (60.60.60.60) (Process ID 1 VRF default)
============================================================

AS External Link State Advertisements (Area 0.0.0.0)
--------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
--------------------------------------------------------------------
40.40.40.40     264     0x80000001  0x24cc4     1           10::/64
40.40.40.40     675     0x80000001  0x5b00f     2           11::/64
```

Showing the LSDB database summary:

```
switch# show ipv6 ospfv3 lsdb database-summary
OSPF Router with ID (10.1.1.1) (Process ID 1 VRF default)
============================================================

Area 0.0.0.0 database summary
--------------------------
LSA Type             Count
--------------------------
Router               2
Network              1
Inter Area Prefix    1
Inter Area Router    0
NSSA External        0
Link                 3
Intra Area Prefix    3
--------------------------
Total                10

Process 1 database summary
--------------------------
LSA Type             Count
--------------------------
Router               2
Network              1
Inter Area Prefix    1
Inter Area Router    0
AS External          2
NSSA External        0
Link                 3
Intra Area Prefix    3
--------------------------
Total                12
```

Showing inter-area prefix LSAs:

```
switch# show ipv6 ospfv3 lsdb inter-area-prefix
OSPF Router with ID (6.6.6.6) (Process ID 1 VRF default)
==========================================================

Inter Area Prefix Link State Advertisements (Area 0.0.0.0)
------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
------------------------------------------------------------------
40.40.40.40     929     0x80000001  0x2498      131072      FEC0:3344::/32
50.50.50.50     928     0x80000001  0x5b2f      65536       111::/64
```

Showing network LSAs:

```
switch# show ipv6 ospfv3 lsdb network
OSPF Router with ID (50.50.50.50) (Process ID 1 VRF default)
==========================================================

Network Link State Advertisements (Area 0.0.0.0)
------------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Router Count
------------------------------------------------------------------------
60.60.60.60     944     0x80000001  0x7179      1360007168  2
50.50.50.50     935     0x80000001  0x516a      19          1
```

Showing NSSA external link states:

```
switch# show ipv6 ospfv3 lsdb nssa-external
OSPF Router with ID (2.2.2.1) (Process ID 1 VRF default)
==========================================================

NSSA External Link State Advertisements (Area 0.0.0.0)
------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
------------------------------------------------------------------
3.3.3.3         264     0x80000001  0x24ac2     1           200::/64
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P	Operator (>) or Manager (#)	

Platforms	Command context	Authority
8360 9300 9300S 10000		

show ipv6 ospfv3 neighbors

```
show ipv6 ospfv3 [<PROCESS-ID>] neighbors [<NEIGHBOR-ID>]
     [interface <INTERFACE-NAME>] [detail | summary]
     [all-vrfs | vrf <VRF-NAME>]
```

Description

Shows information about OSPFv3 neighbors.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 process ID to show OSPFv3 neighbor information for the particular OSPFv3 process. Range: 1 to 65535.
`neighbors` *<NEIGHBOR-ID>*	Shows information about a particular neighbor, specified in IPv4 format (A.B.C.D).
`interface` *<INTERFACE-NAME>*	Shows neighbor information only for the specified interface.
`detail`	Shows detailed information for the neighbors.
`summary`	Shows summary information for the neighbors.
`all-vrfs`	Shows neighbor information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Examples

On the HPE Aruba Networking 6400 Switch Series, interface identification differs.

Showing OSPFv3 neighbors information:

Showing OSPFv3 neighbors information for a specific neighbor:

```
switch# show ipv6 ospfv3 neighbors 3.3.3.3
VRF : default                          Process : 1
------------------------------------------------------------
Router-Id         : 3.3.3.3            Area              : 0.0.0.0
Interface         : 1/1/1              Address           :
fe80::7272:cfff:fe79:7510
State             : FULL               Neighbor Priority : 1
Dead Timer Due    : 00:00:36           Options           : 0x13
Time since last state change : 00h:14m:45s
```

Showing detail OSPFv3 neighbors information for a specific neighbor:

```
switch# show ipv6 ospfv3 neighbors 2.2.2.2 detail
```

```
VRF : default                          Process : 1
-------------------------------------------------------------
Router-Id         : 3.3.3.3            Area              : 0.0.0.0
Interface         : 1/1/1              Address           :
fe80::7272:cfff:fe79:7510
State             : FULL               Neighbor Priority : 1
DR                : 3.3.3.3            BDR               : 1.1.1.3
Dead Timer Due    : 00:00:36           Options           : 0x13
Retransmission Queue Length  : 0
Time Since Last State Change : 00h:14m:45s
```

Showing OSPFv3 neighbors information for interface **1/1/1**:

```
switch# show ipv6 ospfv3 neighbors 3.3.3.3 interface 1/1/1
VRF : default                          Process : 1
-------------------------------------------------------------
Router-Id         : 3.3.3.3            Area              : 0.0.0.0
Interface         : 1/1/1              Address           :
fe80::7272:cfff:fe79:7510
State             : FULL               Neighbor Priority : 1
Dead Timer Due    : 00:00:36           Options           : 0x13
Time Since Last State Change : 00h:14m:45s
```

Showing summary OSPFv3 neighbors information for a specific neighbor for all VRFs:

```
switch# show ipv6 ospfv3 neighbors 3.3.3.3 summary all-vrfs
OSPFv3 Process ID 1 VRF default, Neighbor Summary
=================================================

Interface    Down Attempt Init TwoWay ExStart Exchange Loading Full Total
-------------------------------------------------------------------------
1/1/1        0    0       0    0      0       0        0       1    1
Total        0    0       0    0      0       0        0       1    1

OSPFv3 Process ID 1 VRF red, Neighbor Summary
=============================================

Interface    Down Attempt Init TwoWay ExStart Exchange Loading Full Total
-------------------------------------------------------------------------
1/1/2        0    0       0    0      0       0        0       1    1
Total        0    0       0    0      0       0        0       1    1
```

Showing OSPFv3 neighbors information for VRF red:

```
switch# show ipv6 ospfv3 neighbors vrf red
OSPFv3 Process ID 2 VRF red
===========================

Total Number of Neighbors: 1

Neighbor ID      Priority  State             Interface
-------------------------------------------------------
4.4.4.4          1         FULL/DR           1/1/2
  Neighbor address fe80::7272:cfff:fe79:7510
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ipv6 ospfv3 routes

```
show ipv6 ospfv3 [<PROCESS-ID>] routes [<PREFIX/LENGTH>]
     [all-vrfs | vrf <VRF-NAME>]
```

Description

Shows the OSPFv3 routing table information.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 process ID that shows information from the OSPFv3 routing table for the particular OSPFv3 process. Range: 1 to 65535.
<PREFIX/LENGTH>	Specifies the IPv6 destination prefix showing information about a particular destination prefix. For example, 2010:bd9::/32.
`all-vrfs`	Select to show OSPFv3 routing table information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Examples

On the HPE Aruba Networking 6400 Switch Series, interface identification differs.

Showing OSPFv3 routing table information:

```
switch# show ipv6 ospfv3 routes
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPFv3 Process ID 1 VRF default, Routing Table
----------------------------------------------
```

```
Total Number of OSPFv3 Routes : 2

111::/64           (i) area:0.0.0.0
     directly attached to interface 1/1/1, cost 1 distance 110
fd00::/64          (i) area:0.0.0.1
     directly attached to interface vlan10, cost 1 distance 110
```

Showing OSPFv3 routing table information for VRF red:

```
switch# show ipv6 ospfv3 routes vrf red

Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPFv3 Process ID 2 VRF red, Routing Table
------------------------------------------

Total Number of OSPFv3 Routes : 1

222::/64           (i) area:0.0.0.1
     directly attached to interface 1/1/2, cost 1 distance 110
```

Showing OSPFv3 routing table information for destination prefix fd00::/64:

```
switch# show ipv6 ospfv3 1 routes fd00::/64
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPFv3 Process ID 1 VRF default, Routing Table for prefixes fd00::/64
----------------------------------------------------------------------

Total Number of OSPFv3 Routes : 1

fd00::/64          (i) area:0.0.0.1
     directly attached to interface vlan10, cost 1 distance 110
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P	Operator (>) or Manager (#)	

Platforms	Command context	Authority
8360 9300 9300S 10000		

show ipv6 ospfv3 statistics

```
show ipv6 ospfv3 [<PROCESS-ID>] statistics
     [all-vrfs | vrf <VRF-NAME>]
```

Description

Shows OSPFv3 statistics.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 process ID that shows information on the OSPFv3 SPF statistics for the particular OSPFv3 process. Range: 1 to 65535.
`all-vrfs`	Select to show OSPFv3 SPF statistics information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Examples

Showing OSPFv3 statistics information:

```
switch# show ipv6 ospfv3 statistics
OSPFv3 Process ID 1 VRF default, Statistics (cleared 3h 2m 21s ago)
------------------------------------------------------------------

Unknown Interface Drops         : 0
Unknown Virtual Interface Drops : 0
Bad IPv6 Header Length Drops    : 0
Wrong OSPFv3 Version Drops      : 0
Bad Source IPv6 Drops           : 0
Resource Failure Drops          : 0
Bad Header Length Drops         : 0
Total Drops                     : 0
```

Showing OSPFv3 statistics information for VRF red:

```
switch# show ipv6 ospfv3 2 statistics vrf red

OSPFv3 Process ID 2 VRF red, Statistics (cleared 3h 2m 30s ago)
---------------------------------------------------------------

Unknown Interface Drops         : 0
Unknown Virtual Interface Drops : 0
Bad IPv6 Header Length Drops    : 0
Wrong OSPFv3 Version Drops      : 0
Bad Source IPv6 Drops           : 0
Resource Failure Drops          : 0
Bad Header Length Drops         : 0
```

```
Total Drops                       : 0
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ipv6 ospfv3 statistics interface

```
show ipv6 ospfv3 [<PROCESS-ID>] statistics interface [<INTERFACE-NAME>]
     [all-vrfs | vrf <VRF-NAME>]
```

Description

Shows the OSPFv3 statistics for the OSPFv3-enabled interfaces.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 process ID to show OSPF-enabled interface statistics information on the specified OSPFv3 process. Range: 1 to 65535.
<INTERFACE-NAME>	Selects to show information only for the specified interface.
`all-vrfs`	Select to show OSPF-enabled interface statistics information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Example

Showing OSPFv3-enabled interfaces information for interface 1/1/1:

```
switch# show ipv6 ospfv3 statistics interface all-vrfs
OSPFv3 Process ID 1 VRF default, Interface vlan2000 Statistics  (cleared 0h 2m 52s
```

```
ago)

================================================================================
====

Tx Hello packets      : 0                 Rx Hello packets      : 0
Tx Hello bytes        : 0                 Rx Hello bytes        : 0
Tx DD packets         : 0                 Rx DD packets         : 0
Tx DD bytes           : 0                 Rx DD bytes           : 0
Tx LS request packets : 0                 Rx LS request packets : 0
Tx LS request bytes   : 0                 Rx LS request bytes   : 0
Tx LS update packets  : 0                 Rx LS update packets  : 0
Tx LS update bytes    : 0                 Rx LS update bytes    : 0
Tx LS ack packets     : 0                 Rx LS ack packets     : 0
Tx LS ack bytes       : 0                 Rx LS ack bytes       : 0

Total IPsec packets processed : 0
Total IPsec bytes processed   : 0
Total Number of State Changes : 0
Number of LSAs                : 0
LSA Checksum Sum              : 0

Total OSPFv3 Packets Discarded: 0
----------------------------------

Reason                           Packets Dropped
-------------------------------------------------
Invalid Type                     0
Invalid Length                   0
Invalid Version                  0
Bad or Unknown Source            0
Area Mismatch                    0
Self-originated                  0
Duplicate Router ID              0
Interface Standby                0
Total Hello Packets Dropped      0
  Hello Interval Mismatch        0
  Dead Interval Mismatch         0
  Options Mismatch               0
  MTU Mismatch                   0
  Neighbor Ignored               0
Resource Failures                0
Bad LSA Length                   0
Others                           0
IPsec Authentication Errors      0
IPsec ESP Errors                 0

Total LSAs Ignored : 0
-----------------------
Bad Type           : 0
Bad Length         : 0
Invalid Data       : 0
Invalid Checksum   : 0
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ipv6 ospfv3 virtual-links

```
show ipv6 ospfv3 [<PROCESS-ID>] virtual-links [brief]
   [all-vrfs | vrf <vrf-name>]
```

Description

Displays the current state and parameters of the OSPFv3 virtual links.

Parameter	Description
<PROCESS-ID>	Enter an OSPFv3 process ID to display information on the OSPFv3 virtual links for the particular OSPFv3 process. Range: 1 to 65535.
`brief`	Select to display brief overview information for the OSPFv3 virtual links.
`all-vrfs`	Select to display OSPFv3 virtual links information for all VRFs.
`vrf <VRF-NAME>`	Specifies the name of a VRF. Default: default.

Examples

Show OSPFv3 virtual links information:

```
switch# show ipv6 ospfv3 virtual-links
Virtual link to router 4.4.4.4 is down
---------------------------------------

Process ID 1 VRF default, Transit area 0.0.0.1
Transit delay 1 sec
Timer Intervals: hello 10, dead 40, retransmit 5
Number of Link LSAs: 0, checksum sum 0
0 state changes
AH Authentication: MD5, SPI: 256
```

Show brief overview information for OSPFv3 virtual links:

```
switch# show ospfv3 virtual-links brief
OSPFv3 Process ID 1 VRF default
```

```
================================

Total Number of Virtual Links: 1

Remote Router    Transit Area      Status
-----------------------------------------
4.4.4.4          0.0.0.1           down
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

summary-address

```
summary-address  <IPV6-ADDR>/<MASK> [no-advertise | tag <TAG-VALUE>]
no summary-address <prefix/length> [no-advertise | tag <tag-value>
```

Description

Summarizes the external routes with the matching address and mask. When advertising this route, its metric is set to the lowest cost path from among the routes that were summarized.

The **no** form of this command disables route summarization.

This command only works for an ASBR (Autonomous System Boundary Router).

Parameter	Description
<IPV6-ADDR>	Specifies an IP address in IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F.
<MASK>	Specifies the number of bits in the address mask in CIDR format (**x**), where **x** is a decimal number from 0 to 128.

Parameter	Description
`no-advertise`	Do not advertise the aggregate route. Suppress routes that match the specified prefix/mask pair.
`tag <TAG-VALUE>`	Specify the tag for the aggregate route. The summary prefix will be advertised along with the tag value in External LSAs. Range: 0 to 4294967295

Examples

Setting OSPF route summarization:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# summary-address 2001:DB8::1/32
```

Disabling route summarization:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no summary-address 2001:DB8::1/32
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

timers lsa-arrival

```
timers lsa-arrival <DELAY>
no timers lsa-arrival
```

Description

Configures the minimum delay between receiving the same LSA from a peer. The same LSA is an LSA that contains the same LSA ID number, LSA type, and advertising router ID. If an instance of the same LSA arrives sooner before the delay expires, the LSA is dropped. Generally, the LSA arrival timer should

be set to a value less than or equal to the start-time value for the command **timers throttle lsa start** on the neighbor.

The **no** form of this command sets the LSA timers to default values.

Parameter	Description
<DELAY>	Specifies the delay in milliseconds. Range: 0 to 600000. Default: 1000.

Examples

Setting the LSA arrival timer:

```
switch(config)# router ospf 1
switch(config-ospfv3-1)# timers lsa-arrival 10
```

Setting the LSA arrival timer to default:

```
switch(config)# router ospf 1
switch(config-ospfv3-1)# no timers lsa-arrival
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

timers throttle lsa

```
timers throttle lsa start-time <START-TIME> hold-time <HOLD-TIME> max-wait-time <WAIT-TIME>
no timers throttle lsa
```

Description

Configures the timers for LSA generation.

The **no** form of this command sets the LSA timers to default values.

Parameter	Description
`start-time <START-TIME>`	Specifies the initial wait time in milliseconds after which LSAs are generated. When set to 0, the LSAs are generated without any delay. Range: 0 to 600000. Default: 5000.
`hold-time <HOLD-TIME>`	Specifies the amount of time, in milliseconds, between regeneration of an LSA. The hold time doubles each time the same LSA must be regenerated, until **max-wait-time** is reached. When set to 0, LSA regeneration time is not increased. Range: 0 to 600000. Default: 0.
`max-wait-time <WAIT-TIME>`	Specifies the maximum wait time, in milliseconds, for regeneration of the same LSA. When set to 0, LSA regeneration time is not increased. Range: 0 to 600000. Default: 0.

Examples

Setting the LSA timers:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# timers throttle lsa start-time 100 hold-time 1000 max-
wait-time 10000
```

Setting LSA timers to default values:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no timers throttle lsa
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

timers throttle spf

```
timers throttle spf start-time <START-TIME> hold-time <HOLD-TINME>
max-wait-time <WAIT-TIME>
no timers throttle spf
```

Description

Configures timers for SPF calculation. There are three timers:

- **start-time** Is the initial delay before an SPF calculation is started. Default is 200 milliseconds.
- **hold-time** Is the progressive backoff time to wait before next scheduled SPF calculation. Default is 1000 milliseconds. If a route change event occurs during this period, the value doubles until it reaches the *max-wait-time*.
- **max-wait-time** Is the maximum time to wait before the next scheduled SPF calculation. Default is 5000 milliseconds. This is used to limit the SPF hold timer and also defines the time to be considered for which the OSPF LSDB has to be stable, after which the SPF throttle mechanism is reset.

The **no** form of this command sets all the configured non-default timers to default value.

Parameter	Description
<START-TIME>	Time in milliseconds to set timer for initial SPF delay. Default: 200.
<HOLD-TINME>	Time in milliseconds to set the minimum hold time between two consecutive SPF calculations. Default: 1000.
<WAIT-TIME>	Time in milliseconds to set the maximum wait time between two consecutive SPF calculations. Default: 5000.

Examples

Setting non-default timer values for SPF throttling:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# timers throttling spf start-time 500 hold-time 3000 max-
wait-time 9000
Switch(config-ospfv3-1)# show running-config current-context
router ospfv3 1
               area 0.0.0.0
               area 0.0.0.1
               area 0.0.0.2 nssa no-summary
               area 0.0.0.3 stub
```

Setting default timer values for SPF throttling after configuring non-default values:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no timers throttling spf
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

transit-delay

```
transit-delay <DELAY>
no transit-delay
```

Description

Sets the time delay in Link state transmission for virtual links.

The **no** form of this command sets the delay in Link state transmission to the default of 1 second for virtual links.

Parameter	Description
<DELAY>	Specifies the time delay for the transit delay, in seconds. Range: 1 to 3600. Default: 1.

Examples

Setting OSPFv3 virtual links transit delay:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink6)# transit-delay 30
```

Setting OSPFv3 virtual links transit delay to default:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink6)# no transit-delay
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-router-vlink`	Administrators or local user group members with execution rights for this command.

trap-enable

```
trap-enable
no trap-enable
```

Description

Enables the notification of the events to be sent as traps to the SNMP management stations for OSPFv3.

The **no** form of this command disables the notification of the events to be sent as traps to the SNMP management stations for OSPFv3.

Examples

Enabling sending notification of events as traps:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# trap-enable
```

Disabling sending notification of events as traps:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-1)# no trap-enable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300 9300S 10000		

Open Shortest Path First version 3 Address Family (OSPFv3 AF)

OSPFv3 AF extends the existing OSPFv3 protocol by adding support for multiple unicast address families as described in RFC5838. As an extension of OSPFv3, OSPFv3 AF retains all the functionalities, LSA types, and concepts of the original OSPFv3 protocol.

Overview

The characteristics of OSPFv3 AF are:

- Provides a loop-free topology using SPF algorithm.
- Allows hierarchical routing using area 0 (backbone area) as the top of the hierarchy.
- Supports load balancing with equal cost routes for same destination.
- OSPFv3 AF is a classless protocol and allows for a hierarchical design with VLSM (Variable Length Subnet Masking) and route summarization.
- Scales enterprise size network easily with area concept.
- Provides fast convergence with triggered, incremental updates through Link State Advertisements (LSAs).

OSPFv3 AF configurations are supported in global router context, address-family context, and interface context. Some OSPFv3 AF router configurations are supported only in address-family context, while others are supported in both AF router and global router context. OSPFv3 AF can be configured on routed ports, VLAN interfaces, LAG interfaces, and loopback interfaces. All such configurations work in the mentioned interfaces context. OSPFv3 AF mandates the associated interface to be routing interface.

The switch supports congestion control (prioritizing hello packets, inactivity timer reset, and adjacency throttling).

Supported features

- OSPFv3 neighbor adjacency, Hello protocol, multiple areas, Inter-area routing
- AS external routes, stub areas, totally stubby areas, NSSA, ABR, ASBR
- Designated Router/Backup Designated Router
- Point-to-point interfaces/broadcast interfaces
- Virtual Links (Supported only by IPv6 AF)
- Bidirectional Forwarding Detection (BFD) - refer to the *High Availability Guide* for additional information
- Equal-cost multipath
- Null authentication, Simple password authentication, MD5 authentication

- Area range aggregation - Type-3/Type-7
- External route aggregation - Type-5 LSAs
- Graceful Restart - un-planned, restart interval, Graceful Restart Helper
- Stub router advertisement
- SPF throttling
- LSA Throttling
- Configuration of interface parameters such as priority, cost, hello-interval, dead-interval, retransmit-interval, transit-delay, etc.
- Configuration of virtual link parameters such as hello-interval, dead-interval, retransmit-interval, transit-delay, etc.
- Route redistribution
- Passive interfaces
- Multi VRF support with each VRF having up to eight OSPF process instances
- Congestion control (prioritizing hello packets, inactivity timer reset, and adjacency throttling)

An OSPFv3 AF session will not come up in 6in6 tunnel.

Limitations

- The following configurations are not supported in router context:
  - area
  - bfd all-interfaces
  - default-metric
  - passive-interface-default
  - redistribute
  - maximum-paths
  - traps
  - default-information originate
  - summary-address
  - distribute-list
- Only one instance per address-family can be configured on the interface.
- Multicast address-families are not supported.
- OSPFv3 AF processes cannot establish neighborships with peers that do not support OSPFv3 AF or OSPFv3 Instance-IDs.
- OSPFv3 AF area-level IPsec authentication is not supported.
- Instance ID mismatch counter is not supported per interface. To identify a mismatch in instance IDs, the user must review the configuration of that particular interface. To help identify debug logs for mismatched instance IDs, the user must enable debug ospfv3-af and monitor debug logs for the specific interface.
- The Options field in the show ospfv3-af neighbor detail output does not display the address-family information.
- Simultaneous configurations of OSPFv3-AF with OSPFv2 or OSPFv3 on the same interface are not supported.

How OSPFv3 AF protocol works

OSPFv3 AF protocol

The protocol uses Link State Advertisements (LSAs) transmitted by each router to update neighboring routers regarding its interfaces and the routes available through those interfaces. Each routing switch in an area also maintains a link-state database (LSDB) that describes the area topology. (All routers in a given OSPF area have identical LSDBs). The routing switches used to connect areas to each other flood summary link LSAs and external link LSAs to neighboring OSPF areas to update them regarding available routes. Through this means, each OSPF router determines the shortest path between itself and a desired destination router in the same OSPF domain (Autonomous System (AS)). OSPFv3 AF supports unicast address-families by mapping each AF to a separate Instance ID and OSPFv3 AF instance.

The following ranges are defined for different AFs. The first value of each range is the default value for the corresponding AF.

Instance ID	Description
0 - 31	IPv6 unicast AF
64 - 95	IPv4 unicast AF

OSPFv3 AF concepts

OSPFv3 Address Family is a link-state routing protocol applied to routers grouped into OSPFv3 AF areas identified by the IPv4 and IPv6 routing configuration on each routing switch. Each OSPFv3 AF area includes one or more networks. OSPFv3 AF routers use hello packets and LSAs to maintain OSPFv3 AF operation across networks within an area and between areas within an OSPFv3 AF domain.

OSPFv3 AF uses hello packets to establish and maintain connections between neighboring routers on the same interface. Because neighborship and LSA exchanges are conducted using IPv6 link-local addresses, OSPFv3 AF requires an IPv6 address, even when operating with the IPv4 unicast address family.

OSPFv3 AF Link-state advertisement (LSA) types

OSPFv3 AF uses LSAs transmitted by each router to update neighboring routers regarding its interfaces and the routes available through those interfaces.

Each routing switch in an area also maintains an LSDB that describes the area topology. (All routers in a given OSPFv3 AF area have identical LSDBs, and each router uses the LSDB to build its own shortest-path tree.)

The routing switches used to connect areas to each other flood inter-area-prefix-LSAs, inter-area-router-LSAs, and AS-external-LSAs to backbone area to update it regarding available routes. Each OSPFv3 AF router determines the shortest path between itself and a desired destination router in the same OSPFv3 AF domain (AS). Routed traffic in an OSPFv3 AF AS is classified as one of the following:

- Intra-area traffic
- Inter-area traffic
- External traffic

The routing switches support the LSAs listed in the following table.

Link-state type	Description	Use	Flood scope
0x2001	Router-LSA	Describes the state of each active interface on a router for Area a given area. (Excludes loopback interfaces and interfaces that have not achieved full adjacency.)	Area
0x2002	Network-LSA	Describes the OSPFv3 AF routers in a given network.	Area
0x2003	Inter-area-prefix-LSA	Describes the route to a prefix in another OSPFv3 AF area of Area the same AS. (Excludes prefixes for link-local addresses.) Propagated through backbone area to other areas.	Area
0x2004	Inter-area-router-LSA	Describes the route to an ASBR in another OSPFv3 AF normal area (including the backbone area) of the same AS. Propagated through backbone area to other areas. (Excludes any ASBR in the same area as the router sending the LSA.)	Area
0x4005	AS-external-LSA	Describes the route to a destination prefix in another AS (external route). (Excludes prefixes for link-local addresses.) Originated by ASBR in normal or backbone areas of an AS and propagates through backbone area to other normal areas. Does not flood over virtual links and is not summarized in virtual links. For injection into an NSSA, an NSSA ABR generates a type-7-default-LSA advertising the default route (::/0).	AS
0x2007	NSSA-LSA	Describes the route to a destination in another AS (external route). Originated by ASBR in NSSA. ABR translates type-7 LSAs to AS-external-LSAs for injection into the backbone area.	NSSA

Link-state type	Description	Use	Flood scope
0x0008	Link-LSA	For other routers on the same VLAN interface (except virtual links), describes the router link-local address and any other IPv6 prefixes reachable on the VLAN. Link LSAs are not flooded over virtual links.	Link-local
0x2009	Intra-area-prefix-LSA	Generated on transit links within an area by the DR operating on those links. Also, every OSPFv3 AF router generates this LSA to refer to stub and loopback prefixes on the router.	Area

OSPFv3 AF area types

OSPFv3 AF is built upon a hierarchy of network areas. All areas for a given OSPFv3 AF domain reside in the same AS. An AS is defined as a number of contiguous networks that share an interior gateway routing protocol.

An AS can be divided into multiple areas, including the backbone (area 0). Because each area represents a collection of contiguous networks and hosts, the topology of a given area is not known by the internal routers in any other area. Areas define the boundaries to which router-LSAs and network-LSAs are broadcast. This functionality limits the amount of LSA flooding that occurs within the AS. It also helps to control the size of the link-state databases (LSDBs) maintained in OSPFv3 AF routers.

An area is represented in OSPFv3 AF by either a 32-bit dotted-decimal address or a number. Area types include: Backbone, Normal, Not-so-stubby (NSSA), and Stub.

Backbone area

Every AS must have one (and only one) backbone area (identified as area 0 or 0.0.0.0.) The ABRs of all other areas in the same AS connect to the backbone area, either physically through an ABR or through a configured, virtual link. The backbone is a special type of area and serves as a transit area for carrying the inter-area-prefix-LSAs, AS-external-LSAs, and routed traffic between non-backbone areas, as well as the router-LSAs, network-LSAs, and routed traffic internal to the area. ASBRs are allowed in backbone areas.

Normal area

A normal area allows inter-area-prefix-LSAs and AS-external-LSAs to and from the backbone area. A normal area connects to the AS backbone area through one or more ABRs (physically or through a virtual link) and allows router-LSAs and network LSAs within this area. ASBRs are allowed in normal areas.

Stub area

Stub area connects to the AS backbone through one or more ABRs. It does not allow an internal ASBR, and does not allow AS-external-LSAs. A stub area supports these actions:

- Advertise the area summary routes to the backbone area.
- Advertise inter-area routes from other areas.
- Use the inter-area-prefix-LSA default route to advertise routes to an ASBR and to other areas.

You can configure the stub area ABR to do the following:

- Suppress advertising on the area's summarized internal routes into the backbone area.
- Suppress LSA traffic from other areas in the AS by replacing inter-area-prefix-LSAs and the default external route from the backbone area with the default summary route (::/0.). This area is called totally stubby area.

Virtual links are not allowed for stub areas.

Not-so-stubby (NSSA) area

An NSSA area connects to the backbone area through one or more ABRs. NSSAs are used where an ASBR exists in an area where you want to:

- Block injection of external routes from other areas of the AS.
- Advertise type-7-LSA external routes (learned from the ASBR) to the backbone area as AS-external-LSAs.

NSSAs also support the following:

- Advertise inter-area-prefix-LSAs from the backbone area into the NSSA. (If no-summary is enabled, the NSSA ABR suppresses these LSAs from the backbone and, instead, injects the default route into the NSSA.)
- Advertise NSSA inter-area-prefix-LSAs to the backbone area.

Virtual links are not allowed for NSSAs.

OSPFv3 AF router types

Internal routers

Internal OSPFv3 AF routers belong to only one area. Internal routers flood router-LSAs to all routers in the same area and maintain identical LSDBs.

Area border routers (ABRs)

Area border routers have membership in multiple areas. ABRs are used to connect the various areas in an AS to the backbone area for that AS. Multiple ABRs can be used to connect a given area to the backbone, and a given ABR can belong to multiple areas other than the backbone.

An ABR maintains a separate LSDB for each area to which it belongs. (All routers within the same area have identical LSDBs.) The ABR is responsible for flooding inter-area-prefix-LSAs and inter-area router LSAs between its border areas.

You can reduce summary LSA flooding by configuring area ranges. An area range enables you to assign an aggregate address to a range of IPv6 addresses. This aggregate address is advertised instead of all the individual addresses it represents. You can assign up to eight ranges in an OSPFv3 AF area.

Autonomous system boundary router (ASBR)

Autonomous system boundary routers run one or more interior gateway protocols and serve as a gateway to other autonomous systems operating with interior gateway protocols. The ASBR imports

and translates different protocol routes into OSPFv3 AF through redistribution. ASBRs can be used in backbone areas, normal areas, and NSSAs, but not in stub areas.

Designated routers (DRs)

In an OSPFv3 AF network having two or more routers, one router is elected to serve as the designated router (DR) and another router to act as the backup designated router (BDR). All other routers in the same network segment forward their routing information to the DR and BDR, and the DR forwards this information to all routers in the network. This functionality minimizes the amount of repetitive information that is forwarded on the network by eliminating the need for each individual router in the area to forward its routing information to all other routers in the network. If the area includes multiple networks, each network elects its own DR and BDR.

In an OSPFv3 AF network with no DR and no BDR, the neighboring router with the highest priority is elected the DR and the router with the next highest priority is elected the BDR. If the DR goes off-line, the BDR automatically becomes the DR, and the router with the next highest priority then becomes the new BDR. If multiple routing switches on the same OSPFv3 AF network are declaring themselves DRs, both priority and router ID are used to select the DR and BDR.

Priority is configurable using the `ospfv3-af ipv4 priority` or `ospfv3-af ipv6 priority` command at the interface level. You can use this command to help bias one router as the DR. If two neighbors share the same priority, the router with the highest router ID is designated the DR. The router with the next highest router ID is designated the BDR. Routers retain DR/BDR states throughout the period for which the adjacency is up and running.

OSPFv3 AF configuration task list

Tasks at a glance

- Configuring OSPFv3 AF on the routing switch
- Assigning the routing switch to an OSPFv3 AF area
- Setting OSPFv3 AF network for the area
- Configuring external route redistribution and control
- Configuring area ranges on an ABR to reduce advertisements to the backbone
- Influencing route choice by changing the administrative distance
- Configuring graceful restart
- Configuring OSPFv3 AF virtual link settings
- Configuring OSPFv3 AF interface settings
- Configuring BFD for OSPFv3 AF
- Configuring all OSPFv3 AF interfaces as passive
- Viewing OSPFv3 AF information
- Clearing OSPFv3 AF statistics on a switch

Configuring OSPFv3 AF on the routing switch

Prerequisites

- You must be in the global configuration context, as indicated by the `switch(config)#` prompt to create the OSPFv3 AF instance and enter the OSPFv3 AF configuration context.

- To configure a router ID, and other OSPFv3 AF configuration you must be in the OSPFv3 AF router configuration context, as indicated by the `switch(config-ospfv3-af-1)#` prompt.

Create the OSPFv3 AF instance and enter the OSPFv3 AF configuration context. From this context, you can proceed with other OSPFv3 AF configuration.

Procedure

1. Create the OSPFv3 AF instance and enter the OSPFv3 AF configuration context using the following command. For command details, see router ospfv3-af.

   `router ospfv3-af <PROCESS-ID> [vrf <VRF-NAME>]`

   For example, the following command creates OSPF instance 1.

   ```
   switch(config)# router ospfv3-af 1
   switch(config-ospfv3-af-1)#
   ```

   For example, the following commands enter the IPv4, or IPv6 AF context respectively.

   ```
   switch(config-ospfv3-af-1)# address-family ipv4 unicast
   switch(config-ospfv3-ipv4-uc)#

   switch(config-ospfv3-af-1)# address-family ipv6 unicast
   switch(config-ospfv3-ipv6-uc)#
   ```

2. Configure a global router ID using the following command. For command details, see router-id.

   `router-id <ROUTER-ADDRESS>`

   For example, the following command sets the router ID to 1.1.1.1.

   ```
   switch(config-ospfv3-af-1)# router-id 1.1.1.1
   ```

3. Optionally, if the OSPFv3 AF process was disabled (it is enabled by default), enable the OSPFv3 AF process using the following command.

   `enable`

   For command details, see enable. (Refer to disable for disabling the OSPFv3 AF process).

Assigning the routing switch to an OSPFv3 AF area

Create a Normal, Stub, or Not So Stubby (NSSA) area.

Prerequisites

You must be in the router configuration context, as indicated by the `switch(config-router)#` prompt.

Procedure

Create an OSPFv3 AF area for the routing switch using one of the following commands:

- Create a normal area using the following command: `area <AREA-ID>`. For command details, see area.
- Create a stub area using the following command: `area <AREA-ID> stub`. For command details, see area stub.

- Create a Not-So-Stubby-Area using the command: `area <AREA-ID> nssa`. For command details, see area nssa .

For example, the following commands create either an IPv4 or IPv6 normal area with an area identifier of 1.1.1.1. Area identifier could alternatively be entered in decimal format such as 1.

```
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# area 1
switch(config-ospfv3-ipv4-uc)# area 1.1.1.1
```

```
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 1
switch(config-ospfv3-ipv6-uc)# area 1.1.1.1
```

Setting OSPFv3 AF network for the area

Prerequisites

Routing must be enabled on the interface where you are planning to enable OSPFv3 AF.

You must be in the interface configuration context, as indicated by the `switch(config-if)#` prompt.

After you define an OSPFv3 AF area, you can assign one or more networks to it. The OSPFv3 AF protocol will run on the interface with the configured link-local ipv6 address. The interfaces that have an IPv6 address configured in this network or in a subset of this network will participate in the OSPFv3 AF protocol.

Procedure

1. Set an OSPFv3 AF network for the area using any of the following commands. For command details, see ospfv3-af ipv4 area or ospfv3-af ipv6 area.

   ```
   ospfv3-af <PROCESS-ID> ipv4 area <AREA-ID>
   ospfv3-af <PROCESS-ID> ipv6 area <AREA-ID>
   ```

   For example, use the following command to assign interface 1/1/1 to OSPFv3 AF area 1. The area can alternatively be entered in IO address format.

   ```
   switch(config)# interface 1/1/1
   switch(config-if)# ospfv3-af 1 ipv4 area 1
   ```

   ```
   switch(config)# interface 1/1/1
   switch(config-if)# ospfv3-af 1 ipv6 area 1
   ```

2. Optionally, you can disable OSPFv3 AF on the interface using any of the following commands. For command details, see ospfv3-af ipv4 shutdown or ospfv3-af ipv6 shutdown.

   ```
   ospfv3-af ipv4 shutdown
   ospfv3-af ipv6 shutdown
   ```

   ```
   switch(config)# interface 1
   switch(config-if)# ospfv3-af ipv4 shutdown
   ```

```
switch(config)# interface 1
switch(config-if)# ospfv3-af ipv6 shutdown
```

Configuring external route redistribution and control

Configuring route redistribution for OSPFv3 AF establishes the routing switch as an ASBR for importing and translating different protocol routes from other IGP domains into an OSPFv3 AF domain. When you configure redistribution for OSPFv3 AF, you can specify that routes external to the OSPFv3 AF domain are imported as OSPFv3 AF routes.

1. Enable route redistribution using the `redistribute` command.
   `redistribute {bgp | connected | host-routes | local loopback | static | ripng | ospfv3-af <PROCESS-ID>}`
   `[route-map <ROUTE-MAP-NAME>]`

   For example, setting redistribution of connected routes as OSPFv3 AF routes:

   ```
   switch(config-ospfv3-ipv4-uc)# redistribute connected
   ```

   If a route map is specified, then only the routes that pass the match clause specified in the route map are redistributed to OSPFv3 AF. For example, setting redistribution of local loopback routes, that only match the routes specified by the route map:

   ```
   switch(config-ospfv3-ipv4-uc)# redistribute local loopback route-map local_
   routes
   ```

2. Optionally, modify the default metric for redistribution using the `default-metric` command.
   `default-metric <METRIC-VALUE>`

   For example, setting the default metric for redistribution to 37.

   ```
   switch(config-ospfv3-ipv4-uc)# default-metric 37
   ```

3. Optionally, set the cost of default-summary LSAs using the area command.
   `area <AREA-ID> default-metric <COST>`

   For example, setting the cost of default summary LSAs to 2.

   ```
   switch(config-ospfv3-ipv4-uc)# area 1 default-metric 2
   ```

4. Optionally, use the `max-metric router-lsa` command to set the protocol to advertise a maximum metric so that other routers do not prefer this router as an intermediate hop in their shortest path first (SPF) calculations.
   `max-metric router-lsa [include-stub] | [on-startup <INTERVAL>]`

   For example, setting advertise max-metric router-lsa on startup.

   ```
   switch(config-ospfv3-af-1)# max-metric router-lsa on-startup 3000
   ```

5. Optionally set the maximum number of ECMP routes that OSPFv3 can support using the `maximum-paths` command.

`maximum-paths <MAXIMUM>`

For example, setting the maximum number of ECMP routes to 8.

```
switch(config-ospfv3-ipv4-uc)# maximum-paths 8
```

Configuring area ranges on an ABR to reduce advertisements to the backbone

You can configure area ranges to reduce inter-area advertisements by summarizing a range of IP addresses into a single route advertisement. This action prevents an ABR from advertising specific networks or subnets to the backbone area.

Prerequisites

You must be in the address-family configuration context, as indicated by the `switch(config-ospfv3-ipv4-uc)#` prompt.

Procedure

Summarize inter-area or NSSA paths using the following command. For command details, see area range.

`area <AREA-ID> range <IP-PREFIX> type {inter-area | nssa} [no-advertise]`

For example, use any of the following commands to summarize routes matching the area range fd00::64 using inter-area as the type of address aggregation.

```
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# area 1 range fd00::/64 type inter-area
```

```
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 1 range fd00::/64 type inter-area
```

In another example, use the following command to specify DoNotAdvertise status for routes matching the area range 10.1.0.0./16. Use nssa as the type of address aggregation.

```
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# area 1 range 10.1.0.0/16 type inter-area no
advertise
```

Influencing route choice by changing the administrative distance

Prerequisites

You must be in the OSPFv3 AF router configuration context, as indicated by the `switch(config-ospfv3-af-1)#`prompt.

The administrative distance is used to influence the selection of routes learned by different protocols.

Procedure

Reconfigure the administrative distance using the following command. For command details, see distance.

```
distance [<DISTANCE-VAL> | intra-area [<DISTANCE-VAL>] | inter-area [<DISTANCE-VAL>] |
external [<DISTANCE-VAL>]]
```

For example, use the following command to set administrative distance to 100.

```
switch(config-ospfv3-ipv4-uc)# distance 100
```

Configuring graceful restart

Prerequisites

You must be in the OSPFv3 address-family configuration context, as indicated by the `switch(config-ospfv3-ipv4-uc)#` prompt.

OSPFv3 AF routing can be gracefully restarted on switches without losing packets that are in transit. There is no effect on the saved switch configuration.

Manually create a router-id configuration to support graceful restart and management module failover. A dynamic router-id is chosen from the list of configured IPs at the time of the router-id election, and in the event of a graceful restart and management module failover, the router-id should be preserved before and after these events. Graceful restart would fail if there is a change in the router-id, causing traffic loss.

Procedure

Configure graceful restart of using the following command. For command details, see graceful-restart.

```
graceful-restart {restart-interval <INTERVAL> | helper [strict-lsa-check]}
```

For example, the following command specifies 50 seconds as the maximum interval another router will wait for this router to gracefully restart.

```
switch(config-ospfv3-ipv4-uc)# graceful-restart restart-interval 50
```

Configuring OSPFv3 AF virtual link settings

You must be in the router vlink configuration context, as indicated by the `switch(config-router-vlink)#` prompt.

The OSPFv3 AF interface parameters for this process are automatically set to their default values for virtual links. You can optionally adjust the following OSPFv3 AF virtual link settings.

1. Set the time interval between OSPF hello packets for the OSPFv3 AF virtual links using the following command. For command details, see hello-interval.

   ```
   hello-interval <seconds>
   ```

   For example:

   ```
   switch(config)# router ospfv3-af 1
   switch(config)# address-family ipv6 unicast
   switch(config-ospfv3-ipv4-uc)# area 100 virtual-link 100.0.1.1
   switch(config-router-vlink)# hello-interval 30
   ```

2. Set the interval after which a neighbor is declared dead if no hello packet is received on the OSPFv3 AF virtual links. Use the following command. For command details, see dead-interval.
   `dead-interval <seconds>`
3. Set the time between retransmitting lost link state advertisements for the OSPFv3 AF virtual links using the following command. For command details, see retransmit-interval.
   `retransmit-interval <seconds>`
4. Sets the transit delay in Link state transmission for the OSPFv3 AF virtual links using the following command. For command details, see transit-delay.
   `transit-delay <seconds>`

Configuring OSPFv3 AF interface settings

Prerequisites

You must be in the interface configuration context, as indicated by the `switch(config-if)#` prompt.

You can optionally adjust the following OSPFv3 AF interface settings.

1. Set the interface cost using any of the following commands. For command details, see ospfv3-af ipv4 cost or ospfv3-af ipv6 cost.
   ```
   ospfv3-af ipv4 cost <INTERFACE-COST>
   ospfv3-af ipv6 cost <INTERFACE-COST>
   ```
   For example, the following command sets the cost associated with interface 1/1/1 to 100.

   ```
   switch(config)# interface 1/1/1
   switch(config-if)# ospfv3-af ipv4 cost 100
   ```

   ```
   switch(config)# interface 1/1/1
   switch(config-if)# ospfv3-af ipv6 cost 100
   ```

2. Set the time interval between OSPFv3 AF hello packets for the OSPFv3 AF interface using the following command. For command details, see ospfv3-af ipv6 hello-interval.
   `ospfv3-af ipv6 hello-interval <INTERVAL>`
3. Set the interval after which a neighbor is declared dead if no hello packet is received on the OSPFv3 AF interface using any of the following command. For command details, see ospfv3-af ipv4 dead-interval or ospfv3-af ipv6 dead-interval.
   ```
   ospfv3-af ipv4 dead-interval <INTERVAL>
   ospfv3-af ipv6 dead-interval <INTERVAL>
   ```
4. Set the time between re-transmitting lost link state advertisements for the OSPFv3 AF interface using any of the following command. For command details, see ospfv3-af ipv4 retransmit-interval or ospfv3-af ipv6 retransmit-interval.
   ```
   ospfv3-af ipv4 retransmit-interval <INTERVAL>
   ospfv3-af ipv6 retransmit-interval <INTERVAL>
   ```
5. Sets the time delay in Link state transmission for the OSPFv3 AF interface using any of the following command. For command details, seeospfv3-af ipv4 transit-delay or ospfv3-af ipv6 transit-delay.
   ```
   ospfv3-af ipv4 transit-delay <INTERVAL>
   ospfv3-af ipv6 transit-delay <INTERVAL>
   ```
6. Set the OSPFv3 AF network type for the interface. For command details, see ospfv3-af ipv4 network or ospfv3-af ipv6 network.

```
ospfv3-af ipv4 network {broadcast|point-to-point}
ospfv3-af ipv6 network {broadcast|point-to-point}
```

7. Set the OSPFv3 AF priority on the interface using any of the following command. For command details, see ospfv3-af ipv4 priority or ospfv3-af ipv6 priority.

   ```
   ospfv3-af ipv4 priority <number-value>
   ospfv3-af ipv6 priority <number-value>
   ```

8. Set the OSPFv3 AF interface as OSPFv3 AF passive interface. With this setting the interface participates in OSPFv3 AF but does not send or receive packets on that interface. For command details, see ospfv3-af ipv4 passive or ospfv3-af ipv6 passive.

   ```
   ospfv3-af ipv4 passive
   ospfv3-af ipv6 passive
   ```

Configuring BFD for OSPFv3 AF

Enabling BFD Support for OSPFv3 AF over IPv6 enables OSPFv3 AF to register with Bidirectional Forwarding Detection (BFD) to receive forwarding path detection failure messages. You can either configure BFD for OSPFv3 AF globally on all interfaces or configure it selectively on one or more interfaces. BFD creates a session in asynchronous mode as soon as the switch reaches the 2-Way state with a neighbor. Once the session is established BFD will commence sending echos for path failure detection.

OSPFv3 AF uses IPv6 for neighborship in both IPv4 and IPv6 unicast address-families. Therefore, only a single BFD session is created to monitor the IPv6 link-local address, even when both address-families are enabled on the same interface.

There are two methods for enabling BFD for OSPFv3 AF:

1. Enable BFD for all interfaces enabled for OSPFv3 AF by using the BFD all-interfaces command in address-family configuration mode
2. Enable BFD for a subset of interfaces that have OSPFv3 AF enabled by using the `ospfv3-af bfd` command in interface configuration mode

OSPFv3 AF needs to be enabled prior to enabling BFD on one or more interfaces.

Examples

Enable BFD on all OSPFv3 AF Interfaces

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# bfd all-interfaces
```

Disable BFD on all OSPFv3 AF Interfaces

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# no bfd all-interfaces
```

Enable BFD on a single OSPFv3 AF interface

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af ipv4 bfd
switch(config-if)# ospfv3-af ipv6 bfd
```

Disable BFD on a single OSPFv3 AF interface

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af ipv4 bfd disable
switch(config-if)# ospfv3-af ipv6 bfd disable
```

Set BFD as default on an OSPFv3 AF interface

```
switch(config)# interface 1/1/1
switch(config-if)# no ospfv3-af ipv4 bfd
switch(config-if)# no ospfv3-af ipv6 bfd
```

Configuring all OSPFv3 AF interfaces as passive

Prerequisites

You must be in the OSPFv3 address-family configuration context, as indicated by the `switch(config-ospfv3-ipv4-uc)#` prompt.

OSPFv3 AF sends LSAs to all other routers in the same AS. To limit the flooding of LSAs throughout the AS, you can configure OSPFv3 AF to be passive.

Procedure

Configure all OSPFv3 AF interfaces as passive using the following command. For command details, see passive-interface default.
`passive-interface default`

```
switch(config-ospfv3-ipv4-uc)# passive-interface default
switch(config-ospfv3-ipv6-uc)# passive-interface default
```

Configuring SPF throttling timers

SPF calculation is throttled with default timer values (start-time 200ms, hold-time 1000ms, max-wait-time 5000ms). You can throttle SPF calculation by configuring non-default timers to improve performance of a specific network configuration.

Prerequisites

SPF throttling timers can be configured on any of the following contexts:

- Address-family configuration context, as indicated by the `switch(config-ospfv3-ipv4-uc)#` or `switch(config-ospfv3-ipv6-uc)#` prompt.
- Router configuration context, as indicated by the `switch(config-router)#` prompt.

Procedure

Reconfigure SPF throttling timers using the following command. For command details, see timers throttle spf.

```
timers throttle spf start-time <milliseconds> hold-time <milliseconds> max-wait-time
<milliseconds>
```

For example, use the following command to set start-time to 500ms, hold-time to 3000ms and max-wait-time to 9000ms:

```
switch(config-ospfv3-ipv4-uc)# timers throttle spf start-time 500 hold-time
3000 max-wait-time 9000
```

Viewing OSPFv3 AF information

Prerequisites

Use these show commands from the Manager context, as indicated by the `switch#` prompt.

Procedure

To view OSPFv3 AF information, use the following commands. For command details and examples, click the links.

- To view general OSPFv3 AF, area, state and configuration information use: show ospfv3-af ipv4 or show ospfv3-af ipv6.
- To view information about OSPFv3 AF interfaces use: show ospfv3-af ipv4 interface or show ospfv3-af ipv6 interface.
- To view information about OSPFv3 AF neighbors use: show ospfv3-af ipv4 neighbors or show ospfv3-af ipv6 neighbors .
- To view OSPFv3 AF routing table information use: show ospfv3-af ipv4 routes or show ospfv3-af ipv6 routes.
- To view OSPFv3 AF statistics for the OSPFv3 AF interfaces use show ospfv3-af ipv4 statistics interface or show ospfv3-af ipv6 statistics interface.

Clearing OSPFv3 AF statistics on a switch

Clear OSPFv3 AF event statics using the following command. For command details, see clear ospfv3-af ipv4 statistics or clear ospfv3-af ipv6 statistics . clear ipv6 ospfv3 statistics

```
clear ospfv3-af [<PROCESS-ID>] <address-family> statistics [interface [interface
<INTERFACE-NAME> | vlan <VLAN-ID> | lag <LAG-ID>] | loopback <LOOPBACK-ID> | tunnel
<TUNNEL-ID>] [all-vrfs | vrf <VRF-NAME>]
```

For example, the following command clears OSPFv3 AF event statistics from interface 1/1/1.

```
switch# clear ospfv3-af ipv4 statistics interface 1/1/1
```

```
switch# clear ospfv3-af ipv6 statistics interface 1/1/1
```

OSPFv3 AF commands

active-backbone stub-default-rout

```
active-backbone stub-default-route
```

Description

This command enables the router to send a default route to stub areas if there is an active loopback link in the backbone area. The configuration is not required if backbone area has neighbors or passive interfaces configured. By default, active backbone detection is disabled.

Parameter	Description
stub-default-route	Advertise default route to stub if loopback in backbone area.

Examples

Enabling the router to send an IPv4 default route to stub areas:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# active-backbone stub-default-route
```

Enabling the router to send an IPv6 default route to stub areas:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# active-backbone stub-default-route
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-ipv4-uc<*PROCESS-ID*> config-ospfv3-ipv6-uc<*PROCESS-ID*>	Administrators or local user group members with execution rights for this command.

address-family unicast

```
address-family { <ipv4> | <ipv6> } unicast
```

Description

This command specifies the IPv4/IPv6, address-family and enters address family configuration mode.

Examples

Specifying an IPv4 address-family:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
```

Specifying an IPv6 address-family:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-af-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

area default-metric

```
area <AREA-ID> default-metric <METRIC>
no area <AREA-ID> default-metric
```

Description

Sets the cost of default-summary LSAs announced to the stub/nssa areas.

The **no** form of this command resets the cost of the default-summary LSAs announced to stub/nssa areas to the default of 1.

Parameter	Description
<AREA-ID>	Specifies the area ID as one of the following formats.

Parameter	Description
	OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to 4294967295.
`default-metric <METRIC>`	Specifies the default metric of default-summary LSAs announced to the stub/nssa areas, to the specified value. Default: 1. Range: 0 to 16777215.

Examples

Setting cost for default IPv4 LSA summary:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# area 1 default-metric 2
switch(config-ospfv3-ipv4-uc)# area 1.1.1.1 default-metric 2
```

Setting cost for default IPv4 LSA summary to default:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# no area 1 default-metric
switch(config-ospfv3-ipv4-uc)# no area 0.0.0.1 default-metric
switch(config-ospfv3-ipv4-uc)# no area 0.0.0.1 default-metric 2
```

Setting cost for default IPv6 LSA summary:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 1 default-metric 2
switch(config-ospfv3-ipv6-uc)# area 1.1.1.1 default-metric 2
```

Setting cost for default IPv6 LSA summary to default:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# no area 1 default-metric
switch(config-ospfv3-ipv6-uc)# no area 0.0.0.1 default-metric
switch(config-ospfv3-ipv6-uc)# no area 0.0.0.1 default-metric 2
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-ipv4-uc<PROCESS-ID>` `config-ospfv3-ipv6-uc<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

area nssa

```
area <AREA-ID> nssa [no-summary]
no area <AREA-ID> nssa [no-summary]
```

Description

This command sets the area type as NSSA. If area is present and not NSSA area, changes the area type to NSSA area. If **no-summary** is configured the area type will be NSSA No-Summary.

Parameter	Description
<AREA-ID>	Specifies the area ID is one of the following formats. OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to 4294967295. *Required*
`nssa`	Configure OSPFv3 AF area as NSSA. *Required*
`no-summary`	Do not inject inter-area routes into NSSA. *Optional*
no	Unsets the area type as NSSA. This means that the configured area becomes an area type default. The **no area area-id nssa no-summary** command will stop sending Area Border Router(ABR) summary link advertisements into the NSSA area, but will not unset the area as NSSA. *Optional*

Examples

Creating an NSSA area for OSPFv3 AF:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# area 1 nssa
switch(config-ospfv3-ipv4-uc)# area 1 nssa no-summary
```

Clearing the NSSA area for OSPFv3 AF:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# no area 1 nssa
switch(config-ospfv3-ipv6-uc)# no area 1 nssa no-summary
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-ipv4-uc<PROCESS-ID>` `config-ospfv3-ipv6-uc<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

area

```
area <AREA-ID>
no area <AREA-ID>
```

Description

Creates a normal area with **<AREA-ID>** set if not present. If area is present and is not the normal area, this command changes the area type to normal area.

The **no** form of this command deletes the area with the **<AREA-ID>** specified. The area can be of any type (stub, stub no-summary, and default normal area).

Parameter	Description
<AREA-ID>	Specifies the area ID is one of the following formats. OSPFv3 AF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPFv3 AF area identifier in decimal format. Range: 0 to 4294967295.

Examples

Creating an IPv4 normal area:

```
switch(config)# router ospfv3-af 1
```

```
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# area 1
switch(config-ospfv3-ipv4-uc)# area 1.1.1.1
```

Creating an IPv6 normal area:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 1
switch(config-ospfv3-ipv6-uc)# area 1.1.1.1
```

Deleting an IPv4 area:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# no area 1
switch(config-ospfv3-ipv4-uc)# no area 0.0.0.1
```

Deleting an IPv6 area:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# no area 1
switch(config-ospfv3-ipv6-uc)# no area 1.1.1.1
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-ipv4-uc<*PROCESS-ID*> config-ospfv3-ipv6-uc<*PROCESS-ID*>	Administrators or local user group members with execution rights for this command.

area range

```
area <AREA-ID> range <IP-PREFIX> type {inter-area | nssa} [no-advertise]
no area <AREA-ID> range <IP-PREFIX> type {inter-area | nssa} [no-advertise]
```

Description

Summarizes the routes with the matching address or masks for OSPFv3 AF. This command only works for border routers.

The **no** form of this command unsets the route summarization for the configured IP prefix address on the ABR. When using the **no** form of the command with the **no-advertise** option, enables advertising this range to other areas.

Parameter	Description
<AREA-ID>	Specifies the area identifier as one of the following formats. OSPFv3 AF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPFv3 AF area identifier in decimal format. Range: 0 to 4294967295.
`range` *<IP-PREFIX>*	Specifies summarizing routes matching the area range prefix/mask in IPv4 or IPv6 format.
`type {inter-area | nssa}`	Specifies the type this address aggregation applies to as either inter-area range prefix or NSSA range prefix.
`no-advertise`	Specifies the address range status as **DoNotAdvertise** (do not advertise this range to other areas).

Examples

Summarizing inter-area or NSSA paths on OSPFv3 AF for IPv4:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# area 1 range 10.1.0.0/16 type inter-area
switch(config-ospfv3-ipv4-uc)# area 1 range 10.1.0.0/16 type inter-area no
advertise
switch(config-ospfv3-ipv4-uc)# area 1 range 10.1.0.0/16 type nssa
switch(config-ospfv3-ipv4-uc)# area 1 range 10.1.0.0/16 type nssa no-advertise
```

Unsetting summarization on OSPFv3 AF for IPv4:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# no area 1 range 10.1.0.0/16 type inter-area
switch(config-ospfv3-ipv4-uc)# no area 1 range 10.1.0.0/16 type inter-area no
advertise
switch(config-ospfv3-ipv4-uc)# no area 1 range 10.1.0.0/16 type nssa
switch(config-ospfv3-ipv4-uc)# no area 1 range 10.1.0.0/16 type nssa no-advertise
```

Summarizing inter-area or NSSA paths on OSPFv3 AF for IPv6:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 1 range fd00::1/64 type inter-area
switch(config-ospfv3-ipv6-uc)# area 1 range fd00::1/64 type inter-area no
advertise
switch(config-ospfv3-ipv6-uc)# area 1 range fd00::1/64 type nssa
switch(config-ospfv3-ipv6-uc)# area 1 range fd00::1/64 type nssa no-advertise
```

Unsetting summarization on OSPFv3 AF for IPv6:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# no area 1 range fd00::1/64 type inter-area
switch(config-ospfv3-ipv6-uc)# no area 1 range fd00::1/64 type inter-area no
advertise
switch(config-ospfv3-ipv6-uc)# no area 1 range fd00::1/64 type nssa
switch(config-ospfv3-ipv6-uc)# no area 1 range fd00::1/64 type nssa no-advertise
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-ipv4-uc*<PROCESS-ID>* config-ospfv3-ipv6-uc*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

area stub

```
area <AREA-ID> stub [no-summary]
no area <AREA-ID> stub [no-summary]
```

Description

Creates the stub area with **<AREA-ID>** if not present. If area is present and is not the stub area, this command changes the area type to stub area. If **no-summary** is used, area type will be totally stubby area.

The **no** form of this command unsets the area type as stub. The configured area will be changed to the default normal area. The **no area <AREA-ID> stub no-summary** command will start sending Area Border Router (ABR) summary link advertisements into the stub area, but will not unset the stub area.

ABR does not inject the default route in a Totally Stubby Area with loopback in Area 0.0.0.0. As a workaround, configure a passive interface or active neighbors in the backbone area.

Parameter	Description
<AREA-ID>	Specifies the area ID as one of the following formats. OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to 4294967295.
`stub [no-summary]`	Specifies stub area type. If area is present and not stub area, this parameter changes the area type to stub area. If **no-summary** is specified, area type will be totally stubby area, which means do not inject inter-area routes into stub.

Examples

Creating an IPv4 stub area:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# area 1 stub
switch(config-ospfv3-ipv4-uc)# area 1 stub no-summary
```

Unsetting the IPv4 stub area type:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc) # no area 1 stub
switch(config-ospfv3-ipv4-uc) # no area 1 sub no-summary
```

Creating an IPv6 stub area:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 1 stub
switch(config-ospfv3-ipv6-uc)# area 1 stub no-summary
```

Unsetting the IPv6 stub area type:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc) # no area 1 stub
switch(config-ospfv3-ipv6-uc) # no area 1 sub no-summary
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-ipv4-uc<PROCESS-ID>` `config-ospfv3-ipv6-uc<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

area virtual-link

```
area <AREA-ID> virtual-link <ROUTER-ID>
no area <AREA-ID> virtual-link <ROUTER-ID>
```

Description

Creates an OSPF virtual link with remote ABR (if not created already) and enters the vlink context. The configuration is supported only in ipv6-unicast address-family.

The **no** form of this command deletes an OSPF virtual link with the specified router ID of the remote ABR. If no **<ROUTER-ID>** is specified, the **no** form of the command sets the virtual link to the default settings.

Parameter	Description
<AREA-ID>	Specifies the area ID is one of the following formats. OSPF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPF area identifier in decimal format. Range: 0 to 4294967295.
`virtual-link <ROUTER-ID>`	Configures a virtual link with the specified router ID of the remote ABR.

Examples

Configuring OSPFv3 AF virtual links:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 100 virtual-link 100.0.1.1
```

Deleting OSPFv3 AF virtual links:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# no area 100 virtual-link 100.0.1.1
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-ipv4-uc<PROCESS-ID>` `config-ospfv3-ipv6-uc<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

authentication ipsec spi

```
authentication ipsec spi <SPI-INDEX> <AUTH-TYPE> [<KEY-TYPE> <AUTH-KEY>]
no authentication
```

Description

Configures IPsec AH authentication for the selected Vlink.

The **no** form of this command removes IPsec AH authentication for the selected Vlink.

Parameter	Description
`spi <SPI-INDEX>`	Specifies the Security Parameters Index (SPI) to use. The SPI is an identification tag carried in the IPsec AH header. It enables the receiving OSPF process to select and use the Security Association (SA) from the SA table. The SPI must be unique on the switch. Range: 256 to 4294967295.
<AUTH-TYPE>	Specifies the authentication type: **md5** or **sha1**.
<KEY-TYPE>	Specifies the key type to use: **plaintext** (unencrypted), **hex-string** (encrypted) or **ciphertext** (encrypted).
<AUTH-KEY>	Specifies the authentication key.

When the authentication key is not provided on the command line, plaintext key prompting occurs upon pressing Enter. The entered key characters are masked with asterisks.

Examples

Setting area 1 to use IPsec AH authentication for Vlink with provided plaintext key:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 1
switch(config-ospfv3-ipv6-uc)# area 1 virtual-link 3.3.3.3
switch(config-router-vlink6)# authentication ipsec spi 256 sha1 plaintext abcd
```

Setting area 1 to use IPsec AH authentication for Vlink with prompted plaintext key:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 1
switch(config-ospfv3-ipv6-uc)# area 1 virtual-link 3.3.3.3
switch(config-router-vlink6)# authentication ipsec spi 256 sha1
Enter the IPsec authentication key: **********
Re-Enter the IPsec authentication key: **********
```

Removing IPsec AH authentication for Vlink on area 1:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 1 virtual-link 3.3.3.3
switch(config-router-vlink6)# no authentication
switch(config-router-vlink6)# no authentication ipsec spi 1234 md5 ciphertext
                              qwertyuioplkjhgfdsazxcvbnmkjhgfdsa
switch(config-router-vlink6)# no authentication ipsec spi 1234 sha1 ciphertext
                              qwertyuioplkjhgfdsazxcvbnmkjhgfdsa
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-router-vlink	Administrators or local user group members with execution rights for this command.

clear ospfv3-af ipv4 neighbor

```
clear ospfv3-af [<PROCESS-ID>] ipv4 neighbor [<NEIGHBOR-ID>] [interface <INTERFACE-NAME>
| vlan <VLAN-ID> | lag <LAG-ID>] | loopback <LOOPBACK-ID> | tunnel <TUNNEL-ID>] [all-vrfs
| vrf <VRF-NAME>]
```

Description

Resets the neighbor and clears the OSPFv3 AF neighbor information.

Parameter	Description
<PROCESS-ID>	Specifies the OSPFv3 AF process ID to clear the statistics for the particular OSPFv3 AF process. Range: 1 to 65535.
`neighbor` *<NEIGHBOR-ID>*	Specifies the router ID of a neighbor.
`interface` *<INTERFACE-NAME>*	Specifies the OSPFv3 AF statistics to clear for the specified interface.
`vlan` *<VLAN-ID>*	Specifies the vlan interface name.
`lag` *<LAG-ID>*	Specifies the lag interface name.
`loopback` *<LOOPBACK-ID>*	Specifies the loopback interface name.
`tunnel` *<TUNNEL-ID>*	Specifies the tunnel interface name.
`all-vrfs`	Select to clear the OSPFv3 AF statistics for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF.

Example

Clearing the OSPFv3 AF neighbor information:

```
switch# clear ospfv3-af 1 ipv4 neighbor
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af 1 ipv4 neighbor 3.3.3.3
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af 1 ipv4 neighbor interface 1/1/1
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af 1 ipv4 neighbor 3.3.3.3 vrf red
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af ipv4 neighbor
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af ipv4 neighbor 5.5.5.5
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af ipv4 neighbor interface 1/1/1
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af ipv4 neighbor 5.5.5.5 vrf red
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

clear ospfv3-af ipv4 statistics

```
clear ospfv3-af [<PROCESS-ID>] ipv4 statistics [interface [interface <INTERFACE-NAME> |
vlan <VLAN-ID> | lag <LAG-ID>] | loopback <LOOPBACK-ID> | tunnel <TUNNEL-ID>] [all-vrfs |
vrf <VRF-NAME>]
```

Description

Clears the OSPFv3 AF event statistics.

Parameter	Description
<PROCESS-ID>	Specifies the OSPFv3 AF process ID to clear the statistics for the particular OSPFv3 AF process. Range: 1 to 65535.
`interface` *<INTERFACE-NAME>*	Specifies the OSPFv3 AF statistics to clear for the specified interface.
`vlan` *<VLAN-ID>*	Specifies the vlan interface name.
`lag` *<LAG-ID>*	Specifies the lag interface name.
`loopback` *<LOOPBACK-ID>*	Specifies the loopback interface name.
`tunnel` *<TUNNEL-ID>*	Specifies the tunnel interface name.
`all-vrfs`	Select to clear the OSPFv3 AF statistics for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF.

Example

Clearing the OSPFv3 AF event statistics:

```
switch# clear ospfv3-af ipv4 statistics
switch# clear ospfv3-af ipv4 statistics interface 1/1/1
```

```
switch# clear ospfv3-af ipv4 statistics interface 1/1/1 vrf vrf_red
switch# clear ospfv3-af ipv4 statistics interface tunnel1
switch# clear ospfv3-af ipv4 statistics interface tunnel1 vrf vrf_red
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

clear ospfv3-af ipv6 neighbor

`clear ospfv3-af [<PROCESS-ID>] ipv6 neighbor [<NEIGHBOR-ID>] [interface <INTERFACE-NAME> | vlan <VLAN-ID> | lag <LAG-ID>] | loopback <LOOPBACK-ID> | tunnel <TUNNEL-ID>] [all-vrfs | vrf <VRF-NAME>]`

Description

Resets the neighbor and clears the OSPFv3 AF neighbor information.

Parameter	Description
<PROCESS-ID>	Specifies the OSPFv3 AF process ID to clear the statistics for the particular OSPFv3 AF process. Range: 1 to 65535.
`neighbor` *<NEIGHBOR-ID>*	Specifies the router ID of a neighbor.
`interface` *<INTERFACE-NAME>*	Specifies the OSPFv3 AF statistics to clear for the specified interface.
`vlan` *<VLAN-ID>*	Specifies the vlan interface name.
`lag` *<LAG-ID>*	Specifies the lag interface name.
`loopback` *<LOOPBACK-ID>*	Specifies the loopback interface name.
`tunnel` *<TUNNEL-ID>*	Specifies the tunnel interface name.
`all-vrfs`	Select to clear the OSPFv3 AF statistics for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF.

Example

Clearing the OSPFv3 AF neighbor information:

```
switch# clear ospfv3-af 1 ipv6 neighbor
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af 1 ipv6 neighbor 3.3.3.3
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af 1 ipv6 neighbor interface 1/1/1
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af 1 ipv6 neighbor 3.3.3.3 vrf red
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af ipv6 neighbor
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af ipv6 neighbor 5.5.5.5
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af ipv6 neighbor interface 1/1/1
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
switch# clear ospfv3-af ipv6 neighbor 5.5.5.5 vrf red
Performing clear ospf neighbor may result in traffic disruption.
Do you want to continue (y/n)? y
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

clear ospfv3-af ipv6 statistics

`clear ospfv3-af [<PROCESS-ID>] ipv6 statistics [interface [interface <INTERFACE-NAME> | vlan <VLAN-ID> | lag <LAG-ID>] | loopback <LOOPBACK-ID> | tunnel <TUNNEL-ID>] [all-vrfs | vrf <VRF-NAME>]`

Description

Clears the OSPFv3 AF event statistics.

Parameter	Description
<PROCESS-ID>	Specifies the OSPFv3 AF process ID to clear the statistics for the particular OSPFv3 AF process. Range: 1 to 65535.
`interface` *<INTERFACE-NAME>*	Specifies the OSPFv3 AF statistics to clear for the specified interface.
`vlan` *<VLAN-ID>*	Specifies the vlan interface name.
`lag` *<LAG-ID>*	Specifies the lag interface name.
`loopback` *<LOOPBACK-ID>*	Specifies the loopback interface name.
`tunnel` *<TUNNEL-ID>*	Specifies the tunnel interface name.
`all-vrfs`	Select to clear the OSPFv3 AF statistics for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF.

Example

Clearing the OSPFv3 AF event statistics:

```
switch# clear ospfv3-af ipv6 statistics
switch# clear ospfv3-af ipv6 statistics interface 1/1/1
switch# clear ospfv3-af ipv6 statistics interface 1/1/1 vrf vrf_red
switch# clear ospfv3-af ipv6 statistics interface tunnel1
switch# clear ospfv3-af ipv6 statistics interface tunnel1 vrf vrf_red
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

dead-interval

```
dead-interval <INTERVAL>
no dead-interval
```

Description

Sets the interval after which a neighbor is declared dead if no hello packet comes in for virtual links.

The **no** form of this command sets the dead interval to default for virtual links. The default value is 40 seconds (generally four times the hello packet interval).

Parameter	Description
<INTERVAL>	Specifies the time interval for the dead interval, in seconds. Range: 1 to 65535. Default: 40.

Examples

Setting OSPFv3 AF virtual links dead interval:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink6)# dead-interval 30
```

Setting OSPFv3 AF virtual links dead interval to default:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink6)# no dead-interval
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-router-vlink	Administrators or local user group members with execution rights for this command.

debug ospfv3-af

```
debug ospfv3-af <SUBMODULE-NAME> [instance <PROCESS-ID>] [ip <IP-ADDRESS>] [port <IF-
NAME>] [vrf <VRF-NAME>] [severity {alert | crit | debug | emer | err | info | notice |
warning}]
no debug ospfv3-af <SUBMODULE-NAME> [instance <PROCESS-ID>] [ip <IP-ADDRESS>] [port <IF-
NAME>] [vrf <VRF-NAME>] [severity {alert | crit | debug | emer | err | info | notice |
warning}]
```

Description

Enables problem logs to the debug logs. If <SEVERITY> is omitted, all severities are logged.

The **no** form of this command disables problem logs to the debug logs.

Parameter	Description
<SUBMODULE-NAME>	Enables debug logging for a specific submodule. For a list of supported submodules, enter the debug <MODULE-NAME> command followed by a space and a question mark (?).
`instance` *<PROCESS-ID>*	Displays debug logs for the specified OSPFv3 AF process ID. Range: 1 to 65535.
`ip` *<IP-ADDRESS>*	Displays debug logs for the specified IP Address.
`port` *<INTERFACE-NAME>*	Displays debug logs for the OSPFv3 AF specified interfaces.
`vrf` *<VRF-NAME>*	Displays debug logs for the specified VRF name.
`severity {alert|crit|debug|emer| err|info|notice|warning}`	Enables the debug logs for OSPv3 AF based on the required severity level for the destination. Optional.

Examples

Enables debug logging of event submodules for port 1/1/1 and debug severity:

```
switch# debug ospfv3-af ipv4 event port 1/1/1 severity debug
```

Enables all debug logging for VRF red:

```
switch# debug ospfv3-af ipv4 all vrf vrf_red
```

Disables all OSPFv3 AF debug logging:

```
switch# no debug ospfv3-af ipv4 all
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

default-information originate always

```
default-information originate always [metric <METRIC-VALUE>]
no default-information originate always [metric <METRIC-VALUE>]
```

Description

Configures OSPFv3 AF to advertise the default route (::/0) to its neighbors, regardless if it is present in the routing table or not. Optionally, metric can be set for default route ::/0. The default value is 1.

The **no** form of this command disables advertisement of the default route.

Parameter	Description
metric <METRIC-VALUE	Specifies the OSPF metric value for the default route. Default: 1.

Examples

Setting advertisement of the default route:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# default-information originate always
```

Disabling advertisement of the default route:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# no default-information originate always
```

Setting advertisement of the default route with metric set to 20:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# default-information originate always metric 20
```

Disabling advertisement of the default route and setting the metric to the default value:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# no default-information originate always metric
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-ipv4-uc<*PROCESS-ID*>	Administrators or local user group members with execution rights for this command.

default-information originate

```
default-information originate [metric <METRIC-VALUE>]
no default-information originate [metric <METRIC-VALUE>]
```

Description

Configures OSPFv3 AF to advertise the default route (::/0) to its neighbors if it is present in the routing table. Optionally, the metric value can be set for default route ::/0. The default value is 1.

The **no** form of this command disables advertisement of the default route.

Parameter	Description
metric <METRIC-VALUE>	Specifies the OSPF metric value for the default route. Optional. Default: 1.

Examples

Setting advertisement of the default route and specifying an optional metric value of 20: :

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# default-information originate
switch(config-ospfv3-ipv4-uc)# default-information originate metric 20
```

Disabling advertisement of the default route:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# no default-information originate metric
switch(config-ospfv3-ipv4-uc)# no default-information originate
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-ipv4-uc*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

default-metric

```
default-metric <METRIC-VALUE>
no default-metric
```

Description

Sets the default metric for redistributed routes in the OSPFv3 AF.

The **no** form of this command sets the default metric to be used for redistributed routes into OSPFv3 AF to the default of 25.

Parameter	Description
<METRIC-VALUE>	Specifies the default metric value to use for redistributed routes. Range: 1 to 1677214. Default: 25.

Examples

Setting default metric for redistributed routes:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# default-metric 36
```

Setting default metric for redistributed routes to the default:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1) # address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc) # no default-metric
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-ipv4-uc<*PROCESS-ID*>	Administrators or local user group members with execution rights for this command.

diag ospfv3 lsdb dump

```
diag ospfv3 lsdb dump
```

Description

This command dumps the OSPFv3 AF protocol LSA information.

Examples

```
switch# diag ospfv3 lsdb dump
ospfv3PmAreaLsdbTable row:
ospfv3PmAreaLsdbApplIndex:      537265152
ospfv3PmAreaLsdbAreaId:         0x00000000
ospfv3PmAreaLsdbType:           ospfv3RouterLsa
ospfv3PmAreaLsdbRouterId:       0x01010103
ospfv3PmAreaLsdbLsid:           0x00000000
ospfv3PmAreaLsdbSequence:       -2147483646
ospfv3PmAreaLsdbAge:            872
ospfv3PmAreaLsdbChecksum:       12468
ospfv3PmAreaLsdbAdvertisement:  0x00 0x01 0x20 0x01 0x00 0x00 0x00 0x00
                                0x01 0x01 0x01 0x03 0x80 0x00 0x00 0x02
                                0x30 0xB4 0x00 0x28 0x01 0x00 0x00 0x13
                                0x02 0x00 0x00 0x01 0x10 0x01 0x01 0x01
                                0x10 0x01 0x01 0x01 0x03 0x03 0x03 0x03
ospfv3PmAreaLsdbTypeKnown:      true

ospfv3PmAreaLsdbTable row:
ospfv3PmAreaLsdbApplIndex:      537265152
```

```
ospfv3PmAreaLsdbAreaId:          0x00000000
ospfv3PmAreaLsdbType:            ospfv3RouterLsa
ospfv3PmAreaLsdbRouterId:        0x03030303
ospfv3PmAreaLsdbLsid:            0x00000000
ospfv3PmAreaLsdbSequence:        -2147483646
ospfv3PmAreaLsdbAge:             873
ospfv3PmAreaLsdbChecksum:        64738
ospfv3PmAreaLsdbAdvertisement:   0x00 0x01 0x20 0x01 0x00 0x00 0x00 0x00
                                 0x03 0x03 0x03 0x03 0x80 0x00 0x00 0x02
                                 0xFC 0xE2 0x00 0x28 0x00 0x00 0x00 0x13
                                 0x02 0x00 0x00 0x01 0x10 0x01 0x01 0x01
                                 0x10 0x01 0x01 0x01 0x03 0x03 0x03 0x03
ospfv3PmAreaLsdbTypeKnown:       true

ospfv3PmAreaLsdbTable row:
ospfv3PmAreaLsdbApplIndex:       537265152
ospfv3PmAreaLsdbAreaId:          0x00000000
ospfv3PmAreaLsdbType:            ospfv3NetworkLsa
ospfv3PmAreaLsdbRouterId:        0x03030303
ospfv3PmAreaLsdbLsid:            0x10010101
ospfv3PmAreaLsdbSequence:        -2147483647
ospfv3PmAreaLsdbAge:             878
ospfv3PmAreaLsdbChecksum:        45127
ospfv3PmAreaLsdbAdvertisement:   0x00 0x01 0x20 0x02 0x10 0x01 0x01 0x01
                                 0x03 0x03 0x03 0x03 0x80 0x00 0x00 0x01
                                 0xB0 0x47 0x00 0x20 0x00 0x00 0x00 0x13
                                 0x01 0x01 0x01 0x03 0x03 0x03 0x03 0x03
ospfv3PmAreaLsdbTypeKnown:       true

ospfv3PmAreaLsdbTable row:
ospfv3PmAreaLsdbApplIndex:       537265152
ospfv3PmAreaLsdbAreaId:          0x00000000
ospfv3PmAreaLsdbType:            ospfv3InterAreaPrfxLsa
ospfv3PmAreaLsdbRouterId:        0x01010103
ospfv3PmAreaLsdbLsid:            0x00000001
ospfv3PmAreaLsdbSequence:        -2147483647
ospfv3PmAreaLsdbAge:             913
ospfv3PmAreaLsdbChecksum:        13755
ospfv3PmAreaLsdbAdvertisement:   0x00 0x24 0x20 0x03 0x00 0x00 0x00 0x01
                                 0x01 0x01 0x01 0x03 0x80 0x00 0x00 0x01
                                 0x35 0xBB 0x00 0x24 0x00 0x00 0x00 0x01
                                 0x40 0x00 0x00 0x00 0xFD 0x00 0x00 0x00
                                 0x00 0x00 0x00 0x00
ospfv3PmAreaLsdbTypeKnown:       true

ospfv3PmAreaLsdbTable row:
ospfv3PmAreaLsdbApplIndex:       537265152
ospfv3PmAreaLsdbAreaId:          0x00000000
ospfv3PmAreaLsdbType:            ospfv3IntraAreaPrfxLsa
ospfv3PmAreaLsdbRouterId:        0x03030303
ospfv3PmAreaLsdbLsid:            0x00010000
ospfv3PmAreaLsdbSequence:        -2147483645
ospfv3PmAreaLsdbAge:             873
ospfv3PmAreaLsdbChecksum:        2173
ospfv3PmAreaLsdbAdvertisement:   0x00 0x01 0x20 0x09 0x00 0x01 0x00 0x00
                                 0x03 0x03 0x03 0x03 0x80 0x00 0x00 0x03
                                 0x08 0x7D 0x00 0x2C 0x00 0x01 0x20 0x02
                                 0x10 0x01 0x01 0x01 0x03 0x03 0x03 0x03
                                 0x40 0x00 0x00 0x00 0x01 0x11 0x00 0x00
                                 0x00 0x00 0x00 0x00
ospfv3PmAreaLsdbTypeKnown:       true
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

disable

`disable`

Description

Disables the OSPFv3 AF process. By default OSPFv3 AF process is enabled.

This command does not remove the OSPFv3 AF configurations.

Example

Disabling OSPFv3 AF process:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1) # disable
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325H 8325P 8360 9300 9300S 10000		

distance

```
distance [<DISTANCE-VAL> | intra-area [<DISTANCE-VAL>] | inter-area [<DISTANCE-VAL>] |
external [<DISTANCE-VAL>]]
no distance [<DISTANCE-VAL> | intra-area [<DISTANCE-VAL>] | inter-area [<DISTANCE-VAL>] |
external [<DISTANCE-VAL>]]
```

Description

Defines an administrative distance for OSPFv3 AF. Administrative distance is used as a criteria to select the best route when the same route is learned by multiple routing protocols.

The **no** form of this command sets the OSPFv3 AF administrative distance to the default value of 110. Optionally, administrative distance can be set to default for the specific OSPF route type: intra-area, inter-area, or external type-5 and type-7 routes.

Parameter	Description
<DISTANCE-VAL>	Specifies the OSPFv3 AF administrative distance. Range: 1 to 255. Default: 110.
`intra-area`	Specifies the OSPFv3 AF distance for intra-area routes.
`inter-area`	Specifies the OSPFv3 AF distance for inter-area routes.
`external`	Specifies the OSPFv3 AF distance for external type 5 and type 7 routes.

Usage

Within a given OSPF process, intra-area routes are always given precedence even when distances are configured for inter-area or external type routes.

Examples

Setting OSPFv3 AF administrative distance for address-family ipv6-unicast:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# distance 100
switch(config-ospfv3-ipv6-uc)# distance intra-area 24 external 55 inter-area 66
switch(config-ospfv3-ipv6-uc)# distance intra-area 24 external 55
switch(config-ospfv3-ipv6-uc)# distance external 55
switch(config-ospfv3-ipv6-uc)#exit
```

Setting OSPFv3 AF administrative distance for address-family ipv4-unicast to default:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# no distance
switch(config-ospfv3-ipv4-uc)# no distance external
switch(config-ospfv3-ipv4-uc)# no distance intra-area
switch(config-ospfv3-ipv4-uc)# no distance inter-area
switch(config-ospfv3-ipv4-uc)# no distance 1
switch(config-ospfv3-ipv4-uc)#exit
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-ipv4-uc<PROCESS-ID>` `config-ospfv3-ipv6-uc<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

distribute-list prefix

```
distribute-list prefix <PREFIX-LIST-NAME> {in | out}
no distribute-list prefix <PREFIX-LIST-NAME> {in | out}
```

Description

This command uses an existing prefix list to filter routes that are being installed in the routing table or redistributed to another routing protocol.

The **distribute-list prefix** command filters routes in the inbound or the outbound direction. When this command is issued with the **in** parameter, it filters routes from being installed in the routing table, it does not filter LSAs. When this command is issued with the **out** parameter, it filters only the desired redistributed routes from other protocols.

Parameter	Description
`prefix <PREFIX-LIST-NAME>`	Specify the name of an existing prefix.
`{in | out}`	Select one of the following parameters to set the filter direction: ■ **in**: Filter incoming routes into the routing table ■ **out**: Filter outgoing routing updates

Examples

Enabling the filtering of OSPFv3 AF routes for an IPv4 network, preventing routes from being installed in the routing table or redistributed from another routing protocol.

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# distribute-list prefix list1 in
switch(config-ospfv3-ipv4-uc)# distribute-list prefix list2 out
```

Disabling the filtering of OSPFv3 AF routes for an IPv4 network, allowing routes to be installed in the routing table or redistributed from another routing protocol.

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# no distribute-list prefix list1 in
switch(config-ospfv3-ipv4-uc)# no distribute-list prefix list2 out
```

Enabling the filtering of OSPFv3 AF routes for an IPv6 network, preventing routes from being installed in the routing table or redistributed from another routing protocol.

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# distribute-list prefix list1 in
switch(config-ospfv3-ipv6-uc)# distribute-list prefix list2 out
```

Disabling the filtering of OSPFv3 AF routes for an IPv6 network, allowing routes to be installed in the routing table or redistributed from another routing protocol.

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# no distribute-list prefix list1 in
switch(config-ospfv3-ipv6-uc)# no distribute-list prefix list2 out
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P	config-ospfv3-ipv4-uc<*PROCESS-ID*> config-ospfv3-ipv6-uc<*PROCESS-ID*>	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8360 9300 9300S 10000		

enable

```
enable
```

Description

Enables OSPFv3 AF process when disabled. By default OSPFv3 AF process is enabled.

Example

Enabling OSPFv3 AF process:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# enable
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-af-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

encryption ipsec

```
encryption ipsec spi <SPI-INDEX> <AUTH-TYPE> [<KEY-TYPE> <AUTH-KEY>
   {<ENCR-TYPE> [<KEY-TYPE> <ENCR-KEY>] | null}]
no encryption
```

Description

Configures IPSec ESP authentication and encryption for the selected Vlink.

The **no** form of this command removes IPSec ESP authentication and encryption for the selected Vlink.

Parameter	Description
spi *<SPI-INDEX>*	Specifies the Security Parameters Index (SPI) to use. The SPI is an identification tag carried in the IPsec AH header. It enables the receiving OSPFv3 AF process to select and use the Security Association (SA) from the SA table. The SPI must be unique on the switch. Range: 256 to 4294967295.
<AUTH-TYPE>	Specifies the authentication type: **md5** or **sha1**.
<KEY-TYPE>	Specifies the key type to use: **plaintext** (unencrypted), **hex-string** (encrypted) or **ciphertext** (encrypted).
<AUTH-KEY>	Specifies the authentication key.
<ENCR-TYPE>	Specifies the encryption type: **des**, **3des**, **aes**, or **null**. **NOTE:** Encryption type **aes** is considered to be AES128, AES192, or AES256 based on key length.
<ENCR-KEY>	Specifies the encryption key.

When the authentication key is not provided on the command line, plaintext authentication key prompting occurs upon pressing Enter, followed by encryption type prompting, and finally plaintext encryption key prompting. The entered key characters are masked with asterisks.

When the authentication key and encryption type are provided on the command line but the encryption key is not provided, plaintext encryption key prompting occurs upon pressing Enter. The entered key characters are masked with asterisks.

Examples

Setting area 1 to use IPSec ESP authentication and encryption for Vlink with provided plaintext keys:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 1
switch(config-ospfv3-ipv6-uc)# area 1 virtual-link 3.3.3.3
switch(config-router-vlink6)# encryption ipsec spi 256 md5 plaintext abcd
                              des plaintext abcdefab
```

Setting area 1 to use IPSec md5 authentication and encryption for Vlink with a prompted plaintext authentication key and encryption type:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 1
switch(config-ospfv3-ipv6-uc)# area 1 virtual-link 3.3.3.3
switch(config-router-vlink6)# encryption ipsec spi 256 md5
Enter the IPsec authentication key: ********
Re-Enter the IPsec authentication key: *******

Enter the IPsec encryption type (3des/aes/des/null)? des
```

```
Enter the IPsec encryption key: ********
Re-Enter the IPsec encryption key: ********
```

Setting area 1 to use IPSec sha1 authentication and encryption for Vlink with a provided plaintext authentication key and prompted encryption type:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 1
switch(config-ospfv3-ipv6-uc)# area 1 virtual-link 3.3.3.3
switch(config-router-vlink6)# encryption ipsec spi 256 sha1 plaintext pass des
  ciphertext  Configure ciphertext key.
  hex-string  Configure hex-string key. Key length must be 16 Hex digits.
  plaintext   Configure plaintext key. Key length must be 8 characters.

switch(config-router-vlink6)# encryption ipsec spi 256 sha1 plaintext pass des
Enter the IPsec encryption key: ********
Re-Enter the IPsec encryption key: ********
```

Setting area 1 to use IPSec ESP authentication for Vlink with a provided plaintext authentication key and null encryption:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-1)# area 1
switch(config-ospfv3-1)# area 1 virtual-link 3.3.3.3
switch(config-router-vlink6)# encryption ipsec spi 256 md5 plaintext abcd null
```

Removing IPSec ESP authentication and encryption for Vlink on area 1:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 1 virtual-link 3.3.3.3
switch(config-router-vlink6)# no encryption
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H	`config-router-vlink`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325P 8360 9300 9300S 10000		

graceful-restart ignore-lost-interface

```
graceful-restart ignore-lost-interface
no graceful-restart ignore-lost-interface
```

Description

Enables the restarting router to ignore lost OSPFv3 AF interfaces during a graceful restart process. It is recommended to enable the ignore-lost-interface parameter on the VSF device configured with OSPF enabled on ROPs to ensure the graceful restart completes successfully, as a VSF switchover or failover can render a ROP non-operational, leading to a Router LSA change published to neighbors.

The **no** form of this commands disables the restarting router to ignore OSPFv3 AF interfaces during a graceful restart process.

Parameter	Description
`ignore-lost-interface`	Enable the restarting router to ignore lost OSPFv3 AF interfaces during a graceful restart process. This setting should be enabled on a high availability system to ensure a graceful restart completes successfully, even if OSPFv3 AF-enabled links fail due to High Availability events like a switchover or failover. **NOTE:** Enabling this setting means that the hitless restart procedures do not strictly follow those defined in *RFC 3623, Graceful OSPF Restart*. **NOTE:** OSPFv3 AF neighbors must disable strict LSA checking. If the local node has fewer OSPFv3 AF interfaces after restarting, then the neighbors that were adjacent on those interfaces will clear up their adjacencies to the restarting node and will send out link state updates to advertise the dropped adjacency. If strict LSA checking is enabled, the restarting router's neighbors will exit helper mode when they receive the updated LSAs and the graceful restart will still fail.

Examples

Enabling the restarting router to ignore lost OSPFv3 AF interfaces during a graceful restart process:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-ipv4-uc)# graceful-restart ignore-lost-interface
```

Disabling the restarting router:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-ipv4-uc)# no graceful-restart ignore-lost-interface
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-ipv4-uc<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

graceful-restart

```
graceful-restart {restart-interval <INTERVAL> | helper [strict-lsa-check]}
no graceful-restart {restart-interval <INTERVAL> | helper [strict-lsa-check]}
```

Description

Configures graceful restart for OSPFv3 AF. By default graceful restart is enabled on the OSPFv3 AF router.

The **no** form of this command sets the restart interval to the default of 120 seconds or disables helper mode depending on the specified parameters.

Parameter	Description
`restart-interval <INTERVAL>`	Specifies the time another router waits for this router to gracefully restart and selects the maximum time to wait in seconds. Range: 5 to 1800. Default: 120.
`helper`	Specifies that the router will participate in the graceful restart of a neighbor router.
`strict-lsa-check`	(Optional). Use with the **helper** parameter to enable strict Link state Advertisement (LSA) checking when acting as a restart helper for a restarting peer. **NOTE:** If strict LSA checking is enabled, the restarting router's neighbors will exit helper mode when they receive the updated LSAs and the graceful restart will still fail.

Examples

Enabling OSPFv3 AF graceful restart:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-ipv4-uc)# graceful-restart helper
switch(config-ospfv3-ipv4-uc)# graceful-restart helper strict-lsa-check
switch(config-ospfv3-ipv4-uc)# graceful-restart restart-interval 40
```

Setting the restart interval to default, and disabling helper mode:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-ipv4-uc)# no graceful-restart
switch(config-ospfv3-ipv4-uc)# no graceful-restart helper
switch(config-ospfv3-ipv4-uc)# no graceful-restart helper strict-lsa-check
switch(config-ospfv3-ipv4-uc)# no graceful-restart restart-interval
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-ipv4-uc*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

hello-interval

```
hello-interval <INTERVAL>
no hello-interval
```

Description

Sets the time interval between OSPFv3 AF hello packets for virtual links.

The **no** form of this command sets the hello interval to the default value of 10 seconds for virtual links.

Parameter	Description
<INTERVAL>	Specifies the time interval for the hello interval, in seconds. Range: 1 to 65535. Default: 10.

Examples

Setting OSPFv3 AF virtual links hello interval:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink6)# hello-interval 30
```

Setting OSPFv3 AF virtual links hello interval to default:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink6)# no hello-interval
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-router-vlink`	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv4 retransmit-interval

```
ospfv3-af ipv4 retransmit-interval <INTERVAL>
no ospfv3-af ipv4 retransmit-interval
```

Description

Sets the time between retransmitting lost link state advertisements for the OSPFv3 AF interface.

The **no** form of this command sets the time between retransmitting lost link state advertisements to the default of 5 seconds for the OSPFv3 AF interface.

Parameter	Description
<INTERVAL>	Specifies the retransmit interval, in seconds. Range: 1 to 3600. Default: 5.

Examples

Setting OSPFv3 AF retransmit interval on the interface:

```
switch(config)# interface 1
switch(config-if)# ospfv3-af ipv4 retransmit-interval 30
```

Setting OSPFv3 AF retransmit interval to the default on the interface:

```
switch(config)# interface 1
switch(config-if)# no ospfv3-af ipv4 retransmit-interval
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv4 cost

```
ospfv3-af ipv4 cost <INTERFACE-COST>
no ospfv3-af ipv4 cost
```

Description

Sets the cost (metric) associated with a particular interface. The interface cost is used as a parameter to calculate the best routes.

The **no** form of this command sets the cost (metric) associated with a particular interface to the default of 1. The default cost is calculated automatically depending on the bandwidth of the interface.

Parameter	Description
<INTERFACE-COST>	Specifies the interface cost value. Range: 1 to 65535. Default: 1.

Examples

Setting OSPFv3 AF interface cost:

```
switch(config)# interface 1
switch(config-if)# ospfv3-af ipv4 cost 100
```

Setting the OSPFv3 AF interface cost to default:

```
switch(config)# interface 1
switch(config-if)# no ospfv3-af ipv4 cost
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv4 hello-interval

```
ospfv3-af ipv4 hello-interval <INTERVAL>
no ospfv3-af ipv4 hello-interval
```

Description

Sets the time interval between OSPFv3 AF hello packets for the OSPFv3 AF interface.

The **no** form of this command sets the time interval between OSPFv3 AF hello packets to the default for the OSPFv3 AF interface of 10 seconds.

Parameter	Description
<INTERVAL>	Specifies the time interval between hello packets, in seconds. Range: 1 to 65535. Default: 10.

Examples

Setting OSPFv3 AF hello interval on the interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af ipv4 hello-interval 30
```

Setting OSPFv3 AF hello interval to default on the interface:

```
switch(config)# interface 1/1/1
switch(config-if)# no ospfv3-af ipv4 hello-interval
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if`	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv4 network

```
ospfv3-af ipv4 network {broadcast|point-to-point}
no ospfv3-af ipv4 network {broadcast|point-to-point}
```

Description

Configures the network type for the interface. By default the network type is broadcast network.

The **no** form of this command sets the network type for the interface to the system default as broadcast network for ethernet interface and point-to-point network for tunnel interface.

Parameter	Description
`broadcast`	Specifies the OSPFv3 AF network type as a broadcast multiaccess network.
`point-to-point`	Specifies the OSPFv3 AF network type as a point-to-point network.

Examples

Setting OSPFv3 AF network type for the interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af ipv4 network broadcast
switch(config-if)# ospfv3-af ipv4 network point-to-point
```

Disabling OSPFv3 AF network type for the interface to system default of broadcast network:

```
switch(config)# interface 1/1/1
switch(config-if)# no ospfv3-af ipv4 network
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv4 priority

```
ospfv3-af ipv4 priority <number-value>
no ospfv3-af ipv4 priority
```

Description

Sets the OSPFv3 AF priority for the interface. The larger the numeric value of the priority, the higher the chance it will become the designated router. Setting a priority of 0 makes the router ineligible to become a designated router or back up designated router.

The **no** form of this command sets the OSPFv3 AF priority for the interface to the default of 1.

Parameter	Description
<number-value>	Specifies the priority value. Default: 1. Range: 0 to 255.

Examples

Setting the OSPFv3 AF priority for the interface:

```
switch(config)# interface 1
switch(config-if)# ospfv3-af ipv4 priority 50
```

Setting the OSPFv3 AF priority for the interface to the default of 1:

```
switch(config)# interface 1
switch(config-if)# no ospfv3-af ipv4 priority
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv4 shutdown

```
ospfv3-af ipv4 shutdown
no ospfv3-af ipv4 shutdown
```

Description

Disables OSPFv3 AF on the interface. The interface state changes to Down. It does not remove the interface from the OSPF area. To remove the interface, use the command **no ospfv3-af ipv4 area**.

The **no** form of this command re-enables OSPFv3 AF on the interface.

Examples

Disabling OSPFv3 AF on the interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af ipv4 shutdown
```

Re-enabling OSPFv3 AF on the interface:

```
switch(config)# interface 1/1/1
switch(config-if)# no ospfv3-af ipv4 shutdown
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv4 transit-delay

```
ospfv3-af ipv4 transit-delay <seconds>
no ospfv3-af ipv4 transit-delay <seconds>
```

Description

Sets the time delay in Link state transmission for the OSPFv3 AF interface.

The **no** form of this command sets the transit delay in Link state transmission to the default of 1 second.

Parameter	Description
<DELAY>	Specifies the time delay for the transit delay, in seconds. Range: 1 to 3600. Default: 1.
no	Sets the delay in Link state transmission to the default of 1 for virtual links.

Examples

Setting OSPFv3 AF transit delay on the interface

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af ipv4 transit-delay 30
```

Setting OSPFv3 AF transit delay to default on the interface

```
switch(config)# interface 1/1/1
switch(config-if)# no ospfv3-af ipv4 transit-delay
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv6 area

```
ospfv3-af <PROCESS-ID> ipv6 area <AREA-ID>
no ospfv3-af <PROCESS-ID> ipv6 area <area-id>
```

Description

This command runs the OSPFv3 AF IPv6 on the interface with the configured link-local IPv6 address for the area specified. To move an interface to a new area, unmap the existing area and then associate a new area with the interface. If the OSPFv3 AF process or area is not created, user confirmation is requested to create the OSPFv3 AF process and area.

The **no** form of this command disables OSPFv3 AF on the interface and removes the interface from the area. Interfaces which have an IP address configured on the network or in a subset of the network, stop participating in the OSPF protocol

Parameter	Description
<AREA-ID>	Specifies the area ID is one of the following formats. Area ID in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. Area ID as a decimal value. Range: 0-4294967295.
<PROCESS-ID>	Specifies the OSPFv3 AF process ID. Range: 1 to 65535.

Examples

Setting OSPFv3 AF network for the area:

```
switch(config)# interface 1
switch(config-if)# ospfv3-af 1 ipv6 area 1
switch(config-if)# ospfv3-af 1 ipv6 area 0.0.0.1
```

Disabling OSPFv3 AF network for the area:

```
switch(config)# interface 1
switch(config-if)# no ospfv3-af 1 ipv6 area 1
switch(config-if)# no ospfv3-af 1 ipv6 area 0.0.0.1
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if`	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv6 cost

```
ospfv3-af ipv6 cost <INTERFACE-COST>
no ospfv3-af ipv6 cost
```

Description

Sets the cost (metric) associated with a particular interface. The interface cost is used as a parameter to calculate the best routes.

The **no** form of this command sets the cost (metric) associated with a particular interface to the default which is calculated automatically depending on the bandwidth of the interface.

Parameter	Description
<INTERFACE-COST>	Specifies the interface cost value. Range: 1 to 65535. Default: 1.

Examples

Setting OSPFv3 AF interface cost:

```
switch(config)# interface vlan 1
switch(config-if)# ospfv3-af ipv6 cost 20
```

Setting the OSPFv3 AF interface cost to default:

```
switch(config)# interface vlan 1
switch(config-if)# no ospfv3-af ipv6 cost
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv6 dead-interval

```
ospfv3-af ipv6 dead-interval <seconds>
no ospfv3-af ipv6 dead-interval
```

Description

Sets the interval after a neighbor is declared dead when no hello packet is received on the OSPFv3 AF interface. Dead interval must be greater than hello interval, otherwise, the OSPF protocol will have an undefined behavior.

The **no** form of this command sets the interval after which a neighbor is declared dead, to the default for the OSPFv3 AF interface. The default value is 40 seconds (generally four times the hello packet interval).

Parameter	Description
<seconds>	Specifies the time interval for the dead interval, in seconds. Range: 1 to 65535. Default: 40.

Examples

Setting OSPFv3 AF dead interval on the interface:

```
switch(config)# interface 1
switch(config-if)# ospfv3-af ipv6 dead-interval 30
```

Setting OSPFv3 AF dead interval to default on the interface:

```
switch(config)# interface 1
switch(config-if)# no ospfv3-af ipv6 dead-interval
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv6 hello-interval

```
ospfv3-af ipv6 hello-interval <INTERVAL>
no ospfv3-af ipv6 hello-interval
```

Description

Sets the time interval between OSPFv3 AF hello packets for the OSPFv3 AF interface.

The **no** form of this command sets the time interval between OSPFv3 AF hello packets to the default for the OSPFv3 AF interface of 10 seconds.

Parameter	Description
<INTERVAL>	Specifies the time interval between hello packets, in seconds. Range: 1 to 65535. Default: 10.

Examples

Setting OSPFv3 AF hello interval on the interface:

```
switch(config)# interface 1
switch(config-if)# ospfv3-af ipv6 hello-interval 30
```

Setting OSPFv3 AF hello interval to default on the interface:

```
switch(config)# interface 1
switch(config-if)# no ospfv3-af ipv6 hello-interval
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv6 network

```
ospfv3-af ipv6 network {broadcast|point-to-point}
no ospfv3-af ipv6 network {broadcast|point-to-point}
```

Description

Configures the network type for the interface. By default the network type is broadcast network.

The **no** form of this command sets the network type for the interface to the system default as broadcast network for ethernet interface and point-to-point network for tunnel interface.

Parameter	Description
broadcast	Specifies the OSPFv3 AF network type as a broadcast multiaccess network.
point-to-point	Specifies the OSPFv3 AF network type as a point-to-point network.

Examples

Setting OSPFv3 AF network type for the interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af ipv6 network broadcast
switch(config-if)# ospfv3-af ipv6 network point-to-point
```

Disabling OSPFv3 AF network type for the interface to system default of broadcast network:

```
switch(config)# interface 1/1/1
switch(config-if)# no ospfv3-af ipv6 network
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv6 passive

```
ospfv3-af ipv6 passive
no ospfv3-af ipv6 passive
```

Description

Configures the interface as an OSPFv3 AF passive interface. With this setting, the interface participates in the OSPF, but does not send or receive OSPF packets on that interface.

The **no** form of this command resets the interface as active. With this setting, the interface starts sending and receiving OSPF packets.

Examples

Setting the interface as OSPFv3 AF passive interface:

```
switch(config)# interface vlan 1
switch(config-if)# ospfv3-af ipv6 passive
```

Setting the interface as OSPFv3 AF active interface:

```
switch(config)# interface vlan 1
switch(config-if)# no ospfv3-af ipv6 passive
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv6 priority

```
ospfv3-af ipv6 priority <number-value>
no ospfv3-af ipv6 priority
```

Description

Sets the OSPFv3 AF priority for the interface. The larger the numeric value of the priority, the higher the chance it will become the designated router. Setting a priority of 0 makes the router ineligible to become a designated router or back up designated router.

The **no** form of this command sets the OSPFv3 AF priority for the interface to the default of 1.

Parameter	Description
<number-value>	Specifies the OSPFv3 AF priority value. Default: 1. Range: 0 to 255.

Examples

Setting the OSPFv3 AF priority for the interface:

```
switch(config)# interface vlan /1
switch(config-if)# ospfv3-af ipv6 priority 50
```

Setting the OSPFv3 AF priority for the interface to the default of 1:

```
switch(config)# interface vlan 1
switch(config-if)# no ospfv3-af ipv6 priority
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv6 retransmit-interval

```
ospfv3-af ipv6 retransmit-interval <INTERVAL>
no ospfv3-af ipv6 retransmit-interval
```

Description

Sets the time between retransmitting lost link state advertisements for the OSPFv3 AF interface.

The **no** form of this command sets the time between retransmitting lost link state advertisements to the default 5 seconds.

Parameter	Description
<INTERVAL>	Specifies the time interval for the retransmit interval, in seconds. Range: 1 to 3600. Default: 5

Examples

Setting OSPFv3 AF retransmit interval on the interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af ipv6 retransmit-interval 30
```

Setting OSPFv3 AF retransmit interval to the default on the interface:

```
switch(config)# interface 1/1/1
switch(config-if)# no ospfv3-af ipv6 retransmit-interval
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv6 shutdown

```
ospfv3-af ipv6 shutdown
no ospfv3-af ipv6 shutdown
```

Description

Disables OSPFv3 AF on the interface. The interface state changes to Down. It does not remove the interface from the OSPF area. To remove the interface, use the command **no ospfv3-af ipv6 area**.

The **no** form of this command re-enables OSPFv3 AF on the interface.

Examples

Disabling OSPFv3 AF on the interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af ipv6 shutdown
```

Re-enabling OSPFv3 AF on the interface:

```
switch(config)# interface 1/1/1
switch(config-if)# no ospfv3-af ipv6 shutdown
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if`	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv6 transit-delay

```
ospfv3-af ipv6 transit-delay <seconds>
no ospfv3-af ipv6 transit-delay <seconds>
```

Description

Sets the time delay in Link state transmission for the OSPFv3 AF virtual link.

The **no** form of this command sets the transit delay in Link state transmission to the default of 1 second.

Parameter	Description
<DELAY>	Specifies the time delay for the transit delay, in seconds. Range: 1 to 3600. Default: 1.
no	Sets the delay in Link state transmission to the default of 1 for virtual links.

Examples

Setting OSPFv3 AF transit delay on the interface

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink6)# transit-delay 30
```

Setting OSPFv3 AF transit delay to default on the interface

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
```

```
switch(config-ospfv3-ipv6-uc)# area 100 virtual-link  100.0.1.1
switch(config-router-vlink6)# no transit-delay
switch(config-router-vlink6)# no transit-delay 30
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-router-vlink`	Administrators or local user group members with execution rights for this command.

max-metric router-lsa

```
max-metric router-lsa [include-stub] | [on-startup <INTERVAL>]
no max-metric router-lsa [include-stub] | [on-startup]
```

Description

Sets the protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their shortest path first (SPF) calculations. If the on-startup parameter is used, the router is configured to advertise a maximum metric at start up. That is, for the time specified in seconds, or the default value of 600 seconds. The include-stub parameter configures the router to advertise maximum metric for links including stub links.

To disable advertisement of the maximum metric, use the **no** form of the command.

The **no** form of this command advertises the normal cost metrics instead of advertising the maximized cost metric. This setting causes the router to be considered in traffic forwarding.

Parameter	Description
`on-startup` *<INTERVAL>*	Automatically advertises the stub Router-LSA (or maximize the router-LSA cost metric) for a specified time interval upon OSPFv3 AF startup. Specifies the time in seconds. Range: 5 to 86400. Default: 600.
`include-stub`	Advertise router-LSAs with maximum metric value for stub links.

Examples

Setting to maximize the cost metrics for Router LSA:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# max-metric router-lsa on-startup 3000
```

Setting to advertise the normal cost metrics instead of advertising the maximized cost metric:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# no max-metric router-lsa
switch(config-ospfv3-ipv6-uc)# no max-metric router-lsa on-startup
```

Setting to advertise Router LSA for stub links:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# max-metric router-lsa
switch(config-ospfv3-ipv6-uc)# max-metric router-lsa include-stub on-startup 3000
```

Setting to advertise the normal cost metrics instead of advertising the maximized cost metric for stub links:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-af-1)# no max-metric router-lsa
switch(config-ospfv3-ipv6-uc)# no max-metric router-lsa include-stub on-startup
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-ipv4-uc<PROCESS-ID>` `config-ospfv3-ipv6-uc<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

maximum-paths

```
maximum-paths <MAXIMUM>
no maximum-paths
```

Description

Sets the maximum number of ECMP routes that OSPFv3 AF can support.

The **no** form of this command sets the maximum number of ECMP routes that OSPFv3 AF can support to the default value of 4.

Parameter	Description
<MAXIMUM>	Specifies the maximum number of ECMP routes. Range: 1 to 32. Default: 4.

Examples

Setting maximum number of parallel routes:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)#address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# maximum-paths 3
```

Setting maximum number of parallel paths to the default value of 4:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)#address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc)# no maximum-paths
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-ipv4-uc*<PROCESS-ID>* config-ospfv3-ipv6-uc*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

ospfv3-af encryption ipsec spi

```
ospfv3-af encryption ipsec spi <SPI-INDEX> <AUTH-TYPE>
    [<KEY-TYPE> <AUTH-KEY> <ENCR-TYPE> [<KEY-TYPE> <ENCR-KEY>]]
no ospfv3-af encryption {null| [ipsec spi <SPI-INDEX> <AUTH-TYPE>
    [<KEY-TYPE><AUTH-KEY><ENCR-TYPE> [<KEY-TYPE><ENCR-KEY>]]}
```

Description

Configures IPsec ESP authentication. OSPFv3 interfaces that have IPsec configured at the interface context will not use area level IPsec ESP.

The **no** form of this command removes IPsec ESP for the specified area.

Parameter	Description
`spi <SPI-INDEX>`	Specifies the Security Parameters Index (SPI) to use. The SPI is an identification tag carried in the IPsec AH header. It enables the receiving OSPF process to select and use the Security Association (SA) from the SA table. The SPI must be unique on the switch. Range: 256 to 4294967295.
<AUTH-TYPE>	Specifies the authentication type: **md5** or **sha1**.
<KEY-TYPE>	Specifies the key type to use: **plaintext** (unencrypted), **hex-string** (encrypted) or **ciphertext** (encrypted).
<AUTH-KEY>	Specifies the authentication key.
<ENCR-TYPE>	Specifies the encryption type: **des**, **3des**, **aes**, or **null**. **NOTE:** Encryption type **aes** is considered to be AES128, AES192, or AES256 based on key length.
<ENCR-KEY>	Specifies the encryption key.

When the authentication key is not provided on the command line, plaintext authentication key prompting occurs upon pressing Enter, followed by encryption type prompting, and finally plaintext encryption key prompting. The entered key characters are masked with asterisks.

When the authentication key and encryption type are provided on the command line but the encryption key is not provided, plaintext encryption key prompting occurs upon pressing Enter. The entered key characters are masked with asterisks.

Examples

Setting interface VLAN **1** to use IPsec ESP with provided authentication and encryption keys:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ospfv3-af encryption ipsec spi 256 sha1 plaintext F82
                        des plaintext F82#450b
```

Setting interface VLAN **3** to use IPsec ESP with prompted authentication and encryption keys:

```
switch(config)# interface vlan 3
switch(config-if-vlan)# ospfv3-af encryption ipsec spi 256 sha1
Enter the IPsec authentication key: ********
Re-Enter the IPsec authentication key: ********

Enter the IPsec encryption type (3des/aes/des/null)? des
```

```
Enter the IPsec encryption key: ********
Re-Enter the IPsec encryption key: ********
```

Setting interface VLAN **4** to use IPsec ESP with provided authentication password and encryption type but a prompted encryption key:

```
switch(config)# interface vlan 4
switch(config-if-vlan)# ospfv3-af encryption ipsec spi 256 sha1 plaintext F82 des
Enter the IPsec encryption key: ********
Re-Enter the IPsec encryption key: ********
```

Setting interface VLAN **1** to use IPSec ESP with provided plaintext authentication key and null encryption:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# ospfv3-af encryption ipsec spi 256 sha1 plaintext F82 null
```

Removing IPsec from interface VLAN **1**:

```
switch(config)# interface vlan 1
switch(config-if-vlan)# no ospfv3-af encryption ipsec spi
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ospfv3-af authentication ipsec spi

```
ospfv3-af authentication ipsec spi <SPI-INDEX> <AUTH-TYPE> [<KEY-TYPE> <AUTH-KEY>]
no ospfv3-af authentication ipsec spi
```

Description

Configures IPSec AH authentication. OSPFv3 AF interfaces that have IPsec configured at the interface context will not use area level IPsec. The Security Parameters Index (SPI) is an identification tag carried in IPsec AH header. It enables the receiving OSPF process to select and use the Security Association (SA) from the SA table.

The **no** form of this command removes IPsec AH authentication for the specified area.

IPsec is not supported for 6in6 tunnel interfaces, and IPsec AH SPI must be unique on the router.

Parameter	Description
spi *<SPI-INDEX>*	Specifies the Security Parameters Index (SPI) to use. The SPI is an identification tag carried in the IPsec AH header. It enables the receiving OSPF process to select and use the Security Association (SA) from the SA table. The SPI must be unique on the switch. Range: 256 to 4294967295.
<AUTH-TYPE>	Specifies the authentication type: **md5** or **sha1**.
<KEY-TYPE>	Specifies the key type to use: **plaintext** (unencrypted), **hex-string** (encrypted) or **ciphertext** (encrypted).
<AUTH-KEY>	Specifies the authentication key.

When the authentication key is not provided on the command line, plaintext key prompting occurs upon pressing Enter. The entered key characters are masked with asterisks.

Examples

Setting interface **1/1/1** to use IPsec authentication with a provided plaintext authentication key:

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af authentication ipsec spi 256 md5 plaintext abcd
```

Setting interface **1/1/1** to use IPsec authentication with a prompted plaintext authentication key:

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af authentication ipsec spi 256 md5
Enter the IPsec authentication key: ********
Re-Enter the IPsec authentication key: ********
```

Removing IPsec authentication from interface **1/1/1**:

```
switch(config)# interface 1/1/1
switch(config-if)# no ospfv3-af authentication
switch(config-if)# no ospfv3-af authentication null
switch(config-if)# no ospfv3-af authentication ipsec spi 300 md5 ciphertext
qwertyuioplkjhgfdsxcvbnm
switch(config-if)# no ospfv3-af authentication ipsec spi 300 sha1 ciphertext
qwertyuioplkjhgfdsxcvbnm
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if`	Administrators or local user group members with execution rights for this command.

ospfv3-af encryption ipsec

```
ospfv3-af encryption ipsec spi <SPI-INDEX> <AUTH-TYPE> [<KEY-TYPE> <AUTH-KEY> {<ENCR-TYPE> [<KEY-TYPE> <ENCR-KEY>] | null}]
no ospfv3-af encryption {null | [ipsec spi <SPI-INDEX> <AUTH-TYPE> ciphertext <AUTH-KEY> [<ENCR-TYPE> ciphertext <ENCR-KEY>] | null]}
```

Description

Configures IPsec ESP authentication. OSPFv3 AF interfaces that have IPsec configured at the interface context will not use area level IPsec ESP.

The **no** form of this command removes IPsec ESP for the specified area.

IPsec is not supported in 6in6 tunnel interfaces.

Parameter	Description
`spi <SPI-INDEX>`	Specifies the Security Parameters Index (SPI) to use. The SPI is an identification tag carried in the IPsec AH header. It enables the receiving OSPF process to select and use the Security Association (SA) from the SA table. The SPI must be unique on the switch. Range: 256 to 4294967295.
<AUTH-TYPE>	Specifies the authentication type: **md5** or **sha1**.
<KEY-TYPE>	Specifies the key type to use: **plaintext** (unencrypted), **hex-string** (encrypted) or **ciphertext** (encrypted).
<AUTH-KEY>	Specifies the authentication key.
<ENCR-TYPE>	Specifies the encryption type: **des**, **3des**, **aes**, or **null**. **NOTE:** Encryption type **aes** is considered to be AES128, AES192, or

Parameter	Description
	AES256 based on key length.
<ENCR-KEY>	Specifies the encryption key.

When the authentication key is not provided on the command line, **plaintext** authentication key prompting occurs upon pressing Enter, followed by encryption type prompting, and finally **plaintext** encryption key prompting. The entered key characters are masked with asterisks.

When the authentication key and encryption type are provided on the command line but the encryption key is not provided, **plaintext** encryption key prompting occurs upon pressing Enter. The entered key characters are masked with asterisks.

Examples

Setting interface **1/1/1** to use IPsec ESP with provided authentication and encryption keys:

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af encryption ipsec spi 256 sha1 plaintext abcdef
                        aes plaintext abcdefabcdefabcdefab
```

Setting interface **1/1/1** to use IPSec ESP with provided plaintext authentication key and null encryption:

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af encryption ipsec spi 256 sha1 plaintext abcdef null
```

Setting interface **1/1/1** to use IPsec ESP with prompted authentication and encryption keys:

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af encryption ipsec spi 256 sha1
Enter the IPsec authentication key: **********
Re-Enter the IPsec authentication key: **********

Enter the IPsec encryption type (3des/aes/des/null)? des

Enter the IPsec encryption key: **********
Re-Enter the IPsec encryption key: **********
```

Setting interface **1/1/1** to use IPsec ESP with provided authentication password and encryption type but a prompted encryption key:

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af encryption ipsec spi 256 sha1 plaintext abcdef des
  ciphertext   Configure ciphertext key.
  hex-string   Configure hex-string key. Key length must be 16 Hex digits.
  plaintext    Configure plaintext key. Key length must be 8 characters.
Enter the IPsec encryption key: ********
Re-Enter the IPsec encryption key: ********
```

Removing IPsec from interface **1/1/1**:

```
switch(config)# interface 1/1/1
switch(config-if)# no ospfv3-af encryption
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv4 area

```
ospfv3-af <process-id> ipv4 area <area-id>
no ospfv3-af <process-id> ipv4 area <area-id>
```

Description

This command will run the OSPFv3 AF protocol on the interface with the configured link-local IPv6 address for the area mentioned. For moving an interface to a new area, you need to unmap the existing area and then associate a new area to the interface. If the OSPFv3 AF process or area is not created, user confirmation is requested to create the OSPFv3 AF process and area.

The **no** form of this command disables OSPFv3 AF on the interface.

Parameter	Description
<process-id>	Specify the OSPFv3 AF Process ID. Range: 1-65535
area-id	Area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255 or using a decimal value (0-4294967295).

Examples

Enabling OSPFv3 AF on an interface using IPv4 address:

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af 1 ipv4 area 1
switch(config-if)# ospfv3-af 1 ipv4 area 0.0.0.1
```

Disabling OSPFv3 AF on an interface using IPv4 address:

```
switch(config)# interface 1/1/1
switch(config-if)# no ospfv3-af 1 ipv4 area 1
switch(config-if)# no ospfv3-af 1 ipv4 area 0.0.0.1
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if`	Administrators or local user group members with execution rights for this command.

ospfv3-af ipv4 dead-interval

`ospfv3-af ipv4 dead-interval <seconds>`

Description

This command sets the interval after which a neighbor is declared dead if no hello packet is received for the OSPFv3 AF interface. Dead interval must always be greater than hello-interval, otherwise the OSPFv3 AF protocol will have undefined behavior.

The **no** form of this command sets the interval to the default value of 40 seconds.

Parameter	Description
seconds	Specifies Dead interval value in second. Default is 40. Range: 1-65535

Examples

Configuring an interval of 30 to consider an unresponsive neighbor as down:

```
switch(config)# interface 1
switch(config-if)# ospfv3-af ipv4 dead-interval 30
```

Setting the dead interval to the default value:

```
switch(config)# interface 1
switch(config-if)# no ospfv3-af ipv4 dead-interval
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ospfv3-af authentication null

```
ospfv3-af authentication null
```

Description

This command configures NULL authentication on an interface which disables IPsec AH authentication.

Examples

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af authentication null
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

ospfv3-af encryption null

```
ospfv3-af encryption null
```

Description

This command configures NULL ESP on an interface, which disables IPsec ESP.

Example

Configuring null ESP:

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af encryption null
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

ospfv3-af ipv4 passive

```
ospfv3-af ipv4 passive
```

Description

This command configures the interface as an OSPFv3 AF passive interface. With this setting the interface participates in OSPF but does not receive or send OSPF packets on that interface.

Example

Configuring interface 1/1/1 as an OSPFv3 AF passive interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ospfv3-af ipv4 passive
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

passive-interface default

```
passive-interface default
no passive-interface default
```

Description

Sets all OSPFv3 AF interfaces as passive.

The **no** form of this command sets all the OSPFv3 AF interfaces as active.

Examples

Setting OSPFv3 AF enabled interfaces as passive:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-ipv4-uc)# passive-interface default
```

Setting OSPFv3 AF enabled interfaces as active:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-ipv4-uc)# no passive-interface default
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-ipv4-uc<*PROCESS-ID*>	Administrators or local user group members with execution rights for this command.

redistribute

```
redistribute {bgp | connected | host-routes | local loopback | static | ripng |
   ospf <PROCESS-ID> | ospfv3-af <PROCESS-ID>} [route-map <ROUTE-MAP-NAME>]
no redistribute {bgp | connected | host-routes | local loopback | static | ripng |
   ospf <PROCESS-ID> | ospfv3-af <PROCESS-ID>} [route-map <ROUTE-MAP-NAME>]
```

Description

This command redistributes routes originating from other protocols to OSPFv3 AF. This command allows to redistribute routes from other OSPFv3 AF instances. IPv4 address family allows redistribution from other OSPFv3 Ipv4 AFs. IPv6 address family allows redistribution from other OSPFv3 Ipv6 AFs. In addition, if a route-map is specified, then routes that pass the match clause specified in the route-map are only redistributed to OSPFv3 AF. Configuration is not allowed if the referenced route-map has not yet been configured.

Only BGP leaked routes will be given to OSPFv3 AF for redistribution. If you try to redistribute routes from an OSPF process which is not created, user confirmation is asked to create the process before proceeding with redistribution. If accepted, the process is created with defaults and redistribution configuration is applied. If denied, redistribution configuration is skipped.

If command **route-redistribute active-routes-only** has been issued, only the routes from other protocols which are selected for forwarding are considered for redistribution into OSPFv3 AF. With the change in redistribute behavior, additional explicit redistribution of connected routes may be needed if any of the deployment assumes redistribution of connected networks will happen via redistribution of dynamic protocol routes into OSPF.

The **no** form of this command disables redistribution of routes to the current OSPFv3 AF process.

Parameter	Description
`bgp`	Specifies redistributing BGP (Border Gateway Protocol) routes into OSPFv3 AF.
`connected`	Specifies redistributing connected (directly attached subnet or host).
`host-routes`	Redistribute dynamic, IVRL and static ARP/ND entries into OSPFv3 AF. ARP/ND entries learnt from EVPN will not be redistributed. The **host-routes** parameter is applicable to HPE Aruba Networking 8325, 10000, 8360, and 9300/9300S Switch series only.
`local loopback`	Specifies redistributing local routes of the loopback interface.
`static`	Specifies redistributing static routes into OSPFv3 AF.
`ripng`	Specifies redistributing RIPng routes into OSPFv3 AF.
`ospf <PROCESS-ID>`	Specifies redistributing routes from the specified OSPF process ID. Range: 1 to 65535.
`ospfv3-af <PROCESS-ID>`	Specifies redistributing routes from the specified OSPFv3 AF process ID. Range: 1 to 65535.
`route-map <ROUTE-MAP-NAME>`	Specifies redistribution filtering by route map. To create a route map, use command **route-map**.

Examples

Redistributing routes to OSPFv3 AF:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-ipv4-uc) # redistribute bgp
switch(config-ospfv3-ipv4-uc) # redistribute connected
switch(config-ospfv3-ipv4-uc) # redistribute host-routes
switch(config-ospfv3-ipv4-uc) # redistribute static
switch(config-ospfv3-ipv4-uc) # redistribute ripng
switch(config-ospfv3-ipv4-uc) # redistribute ospfv3-af 2
switch(config-ospfv3-ipv4-uc) # redistribute local loopback
switch(config-ospfv3-ipv4-uc) # redistribute static route-map
% Command incomplete.
switch(config-ospfv3-ipv4-uc) # redistribute static route-map static_networks
switch(config-ospfv3-af-1)#  redistribute static route-map static_networks1
Route Map static_networks1 not found
switch(config-ospfv3-af-1)#  redistribute static route-map static_networks12
Redistribute with route-map config exists
switch(config-ospfv3-ipv4-uc) # redistribute ripng route-map
% Command incomplete.
switch(config-ospfv3-ipv4-uc) # redistribute ripng route-map ripng_networks
switch(config-ospfv3-af-1)#  redistribute ripng route-map ripng_networks1
Route Map ripng_networks1 not found
switch(config-ospfv3-af-1)#  redistribute ripng route-map ripng_networks12
Redistribute with route-map config exists
switch(config-ospfv3-ipv4-uc) # redistribute ospfv3-af 2 route-map
% Command incomplete.
```

```
switch(config-ospfv3-ipv4-uc) # redistribute ospfv3-af 2 route-map tagged_routes
switch(config-ospfv3-ipv4-uc)#  redistribute ospfv3-af 2 route-map tagged_routes1
Route Map tagged_routes1 not found
switch(config-ospfv3-ipv4-uc)#  redistribute ospfv3-af 2 route-map tagged_routes12
Redistribute with route-map config exists
switch(config-ospfv3-ipv4-uc)# redistribute local loopback route-map
% Command incomplete.
switch(config-ospfv3-ipv4-uc)# redistribute local loopback route-map local_routes
switch(config-ospfv3-af-1)# redistribute local loopback route-map local_routes1
Route Map local_routes1 not found
switch(config-ospfv3-ipv4-uc)# redistribute local loopback route-map local_
routes12
Redistribute with route-map config exists
switch(config-ospfv3-ipv4-uc)# redistribute host-routes route-map
% Command incomplete.
switch(config-ospfv3-ipv4-uc)# redistribute host-routes route-map host_routes
switch(config-ospfv3-ipv4-uc)# redistribute host-routes route-map host_routes1
Route Map local_routes1 not found
switch(config-ospfv3-ipv4-uc)# redistribute host-routes route-map host_routes12
Redistribute with route-map config exists
```

Disabling redistributing routes to OSPFv3 AF:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-ipv6-uc) # no redistribute bgp
switch(config-ospfv3-ipv6-uc) # no redistribute connected
switch(config-ospfv3-ipv6-uc) # no redistribute host-routes
switch(config-ospfv3-ipv6-uc) # no redistribute static
switch(config-ospfv3-ipv6-uc) # no redistribute ripng
switch(config-ospfv3-ipv6-uc) # no redistribute ospfv3-af 2
switch(config-ospfv3-ipv6-uc) # no redistribute local loopback
switch(config-ospfv3-ipv6-uc) # no redistribute static route-map
% Command incomplete.
switch(config-ospfv3-ipv6-uc) # no redistribute static route-map static_networks
switch(config-ospfv3-ipv6-uc) # no redistribute ripng route-map
% Command incomplete.
switch(config-ospfv3-ipv6-uc) # no redistribute ripng route-map ripng_networks
switch(config-ospfv3-ipv6-uc) # no redistribute ospfv3-af 2 route-map
% Command incomplete.
switch(config-ospfv3-ipv6-uc) # no redistribute ospfv3-af 2 route-map tagged_
routes
switch(config-ospfv3-ipv6-uc) # no redistribute local loopback route-map
% Command incomplete.
switch(config-ospfv3-ipv6-uc) # no redistribute local loopback route-map local_
routes
switch(config-ospfv3-ipv6-uc)# no redistribute host-routes route-map
% Command incomplete.
switch(config-ospfv3-ipv6-uc)# no redistribute host-routes route-map host_routes
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-ipv4-uc<PROCESS-ID>` `config-ospfv3-ipv6-uc<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

reference-bandwidth

```
reference-bandwidth <BANDWIDTH-VALUE>
no reference-bandwidth
```

Description

Sets the reference bandwidth for OSPFv3 AF. If the OSPFv3 AF interface cost is not explicitly set, then the cost of all the OSPFv3 AF interfaces is recalculated based on the reference bandwidth and link speed of the interface.

For VLAN interfaces the calculated link speed value is 1 Gbps (if the OSPFv3 AF interface cost is not explicitly set).

The **no** form of this command sets the reference bandwidth for OSPFv3 AF to the default of 100000 Mbps.

Parameter	Description
<BANDWIDTH-VALUE>	Specifies the reference bandwidth used to calculate the cost of an interface in Mbps. Range: 1 to 4000000. Default: 100000.

Examples

Setting the reference bandwidth:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-1)# reference-bandwidth 40000
```

Setting the reference bandwidth to the default value:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-1)# no reference-bandwidth
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

retransmit-interval

```
retransmit-interval <INTERVAL>
no retransmit-interval
```

Description

Sets the time between retransmitting lost link state advertisements for virtual links.

The **no** form of this command sets the time between retransmitting lost link state advertisements to the default of 5 seconds for virtual links.

Parameter	Description
<INTERVAL>	Specifies the time interval for the retransmit interval, in seconds. Range: 1 to 3600. Default: 5.

Examples

Setting OSPFv3 AF virtual links retransmit interval:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink6)# retransmit-interval 30
```

Setting OSPFv3 AF virtual links retransmit interval to default:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink6)# no retransmit-interval
switch(config-router-vlink6)# no retransmit-interval 30
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-router-vlink`	Administrators or local user group members with execution rights for this command.

router-id

```
router-id <ROUTER-ADDRESS>
no router-id
```

Description

Sets an ID for the router in an IPv4 address format.

Parameter	Description
<ROUTER-ADDRESS>	Specifies the router address in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. Required
`no`	Unconfigures the router-id for the instance. The Router-id is changed to the dynamically selected router-id. The default router-id 0.0.0.0 is updated to the routing stack that triggers to auto-elect a router-id based on the highest ip-address of loopback interface or highest ip-address of interfaces. The configuration is supported globally as well as per address-families ipv4-unicast, ipv6-unicast. If address-family specific router-id is configured, it takes precedence over global configuration.

Examples

Setting router-id in the OSPFv3 AF context:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1) # router-id 1.1.1.1
switch(config-ospfv3-af-1) # address-family ipv4 unicast
switch(config-ospfv3-ipv4-uc) # router-id 2.2.2.2
switch(config-ospfv3-af-1) # address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc) # router-id 3.3.3.3
```

Unconfiguring router-id:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-ipv4-uc)# no router-id
switch(config-ospfv3-ipv4-uc)# no router-id 1.1.1.1
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-ipv4-uc<PROCESS-ID>` `config-ospfv3-ipv6-uc<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

router ospfv3-af

```
router ospfv3-af <PROCESS-ID> [vrf <VRF-NAME>]
no router ospfv3-af <PROCESS-ID> [vrf <VRF-NAME>]
```

Description

Creates the OSPFv3 AF process (if not created already) and enters the router OSPFv3 AF instance context. Optionally if specified, you can specify a named VRF, or the default VRF if the <*vrf-name*> is not specified. Maximum of eight OSPFv3 processes are allowed per VRF.

The switch allows configuration of OSPFv2, OSPFv3, OSPFv3 IPv4 AF and OSPFv3 IPv6 AF. But at any point only 8 OSPF instances can run per address-family, 8 instances for IPv4 (OSPFv2 + OSPFv3 IPv4 AF) and 8 instances for IPv6 (OSPFv3 + OSPFv3 IPv6 AF).

The **no** form of this command removes the OSPFv3 AF instance. If a VRF is specified, it removes the OSPF instance from the named VRF, or the default VRF if the <*var-name*> is not specified.

Parameter	Description
`<PROCESS-ID>`	Specifies an OSPFv3 AF process ID. Length: 1 to 65535.
`vrf <VRF-NAME>`	Specifies the name of a VRF. Default: default.

Examples

Entering the router OSPFv3 AF instance:

```
switch(config)# router ospfv3-af 1
```

Setting the router OSPFv3 VRF instance:

```
switch(config)# router ospfv3-af 1 vrf vrf_red
```

Removing the router OSPFv3 AF instance:

```
switch#(config)# no router ospfv3-af 1
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

show ospfv3-af ipv4 statistics

`show ospfv3-af [<process-id>] ipv4 statistics [all-vrfs | vrf <vrf-name>]`

Description

This command displays the OSPFv3 AF shortest path first (SPF) calculation statistics.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 AF process ID. Range: 1 to 65535
`all-vrfs`	Display the OSPFv3 AF statistics information for all the VRFs.
`vrf` *<VRF-NAME>*	Specify the VRF name.

Example

```
switch# show ospfv3-af ipv4 statistics
Address-family : IPv4 Unicast
-----------------------------
OSPFv3 AF Process ID 1 VRF default, Statistics (cleared 3h 2m 21s ago)
-----------------------------------------------------------------
Unknown Interface Drops         : 0
Unknown Virtual Interface Drops : 0
Bad IPv6 Header Length Drops    : 0
Wrong OSPFv3 Version Drops      : 0
Bad Source IPv6 Drops           : 0
Resource Failure Drops          : 0
Bad Header Length Drops         : 0
Total Drops                     : 0
switch# show ospfv3-af 2 ipv4 statistics vrf red
Address-family : IPv4 Unicast
-----------------------------
OSPFv3 AF Process ID 2 VRF red, Statistics (cleared 3h 2m 30s ago)
-----------------------------------------------------------------
Unknown Interface Drops         : 0
Unknown Virtual Interface Drops : 0
Bad IPv6 Header Length Drops    : 0
Wrong OSPFv3 Version Drops      : 0
Bad Source IPv6 Drops           : 0
Resource Failure Drops          : 0
Bad Header Length Drops         : 0
Total Drops                     : 0
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show debug buffer module ospfv3-af

```
show debug buffer module ospfv3-af
   [severity (alert|crit|debug|emer|err|info|notice|warn)]
```

Description

Displays debug logs for OSPFv3 AF stored in the specified debug buffer.

Parameter	Description
<MODULE-NAME>	Filters debug logs displayed by the specified module name.
`severity (alert|crit|debug|emer| err|info|notice|warn)`	Displays debug logs for OSPFv3 AF with a specified severity level. Defaults to**debug**. Optional.
`alert`	Displays debug logs with a severity level of **alert (6)** and above.
`crit`	Displays debug logs with a severity level of **critical (5)** and above.
`debug`	Displays debug logs with all severities.
`emer`	Displays debug logs with a severity level of **emergency (7)** only.
`err`	Specifies storage of debug logs with severity level of **error (4)** and above.
`info`	Displays debug logs with a severity level of **info (1)** and above.
`notice`	Specifies storage of debug logs with severity level of **notice (2)** and above.
`warning`	Displays debug logs with a severity level of **warning (3)** and above.

Examples

```
switch(config)# show debug

--------------------------------------------------------------------------------
-------
module      sub_module      severity   vlan  port     ip       mac       instance
     vrf
--------------------------------------------------------------------------------
--------
ospfv3-af   all             debug

switch(config)# show debug buffer module ospfv3-af
--------------------------------------------------------------------------------
-----------------------------------------------------
show debug buffer
--------------------------------------------------------------------------------
-----------------------------------------------------
2017-09-18:10:35:38.621611|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|OSPF
537265152 hitless restart state changed (None -> None),
reason: None.
2017-09-18:10:35:38.621626|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|(OSPFv3
process Id = 1, VRF Name = default)
2017-09-18:10:35:38.621836|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-
AF|0X01111000 Protocol Manager has initialized successfully.
2017-09-18:10:35:38.621846|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|Entity
index              = OSPF 537265152
2017-09-18:10:35:38.621854|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|(OSPFv3
process Id = 1, VRF Name = default)
2017-09-18:10:35:38.622188|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|OSPF
537265152 hitless restart state changed (None -> None),
reason: None.
2017-09-18:10:35:38.622197|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|(OSPFv3
process Id = 1, VRF Name = default)
```

```
2017-09-18:10:35:38.622423|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|OSPF
537265152 created.
2017-09-18:10:35:38.622433|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|(OSPFv3
process Id = 0, VRF Name = default)
2017-09-18:10:35:38.622516|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|A master
join has been successfully registered.
2017-09-18:10:35:38.622524|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|Join
interface type       = 2
2017-09-18:10:35:38.622536|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|Join
interface ID         = 0X00000001
2017-09-18:10:35:38.622544|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|(OSPFv3
process Id = 1, VRF Name = default)
2017-09-18:10:35:38.622717|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|OSPF
537265152 i/f idx 0X10010101: OSPF has been told over
the I3 interface about a new address on an
2017-09-18:10:35:38.622728|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-
AF|interface, or about updated attributes of an existing address.
2017-09-18:10:35:38.622738|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|Address
                        = 111::1
2017-09-18:10:35:38.622746|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|Scope ID
                        = 268501249
2017-09-18:10:35:38.622754|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|Prefix
length                   = 64
2017-09-18:10:35:38.622762|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|Remote
router address           = ::
2017-09-18:10:35:38.622771|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|Address
use flags                = 0X00
2017-09-18:10:35:38.622780|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|(OSPFv3
process Id = 1, VRF Name = default)
2017-09-18:10:35:38.625212|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|OSPF
537265152 activated successfully
2017-09-18:10:35:38.625225|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|Router
ID             rtr ID 1.1.1.1
2017-09-18:10:35:38.625235|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|(OSPFv3
process Id = 1, VRF Name = default)
2017-09-18:10:35:38.625286|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|OSPF
537265152 was Down and is now Going Up.
2017-09-18:10:35:38.625295|hpe-routing|LOG_INFO|AMM|-|OSPFV3-AF|OSPFV3-AF|(OSPFv3
process Id = 1, VRF Name = default)
2017-09-18:10:35:38.625309|hpe-routing|LOG_ERR|AMM|-|OSPFV3-AF|OSPFv3_AF_
EVENT|Failed to configure interface cost for stub link port 1/1/1.
2017-09-18:10:35:38.625312|hpe-routing|LOG_ERR|AMM|-|OSPFV3-AF|OSPFv3_AF_
EVENT|Could not send ospfv3VirtNbrStateChange Trap
2020-07-21:11:36:47.324342|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|OSPF 537265152 NM has sent a Hello packet.
2020-07-21:11:36:47.324717|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|Sending interface index =  i/f idx 0X51101000.
2020-07-21:11:36:47.324952|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|Router ID of target neighbor
(null string if sent per interface) = .
2020-07-21:11:36:47.325217|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|Destination IP address = IP addr FF02::5.
2020-07-21:11:36:47.325460|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_PACKET|
(End of packet, OSPFv3 process Id = 1, VRF Name = default)
2020-07-21:11:36:47.326710|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|OSPF message details:
2020-07-21:11:36:47.326946|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|Header: Src IP:FF02::5 Dst IP:FE80::800:901:4C3:C58B
Type: 1 Rtr Id:2.2.2.2 Area:0.0.0.1 Inst ID:0
2020-07-21:11:36:47.327166|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|Packet: Hello: I/f 0X51101000 Pri:1 Opt:19 Hello
Int:10 Dead Int:40 DR:2.2.2.2 BDR:1.1.1.1
```

```
2020-07-21:11:36:47.327350|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_PACKET|
(End of packet, OSPFv3 process Id = 1, VRF Name = default)
2020-08-12:11:59:50.183653|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|OSPF message details:
2020-08-12:11:59:50.183666|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|Header: Src IP:FE80::800:901:4C1:3751 Dst IP:FF02::5
Type: 4 Rtr Id:2.2.2.2 Area:0.0.0.1 Inst ID:0
2020-08-12:11:59:50.183676|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|Packet: LS Upd: Num LSAs 2
2020-08-12:11:59:50.183687|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_PACKET|
(End of packet, OSPFv3 process Id = 1, VRF Name = default)
2020-08-12:11:59:50.272301|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|OSPF message details:
2020-08-12:11:59:50.272313|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|Header: Src IP:FE80::800:901:4C1:3751 Dst IP:FF02::5
Type: 5 Rtr Id:2.2.2.2 Area:0.0.0.1 Inst ID:0
2020-08-12:11:59:50.272323|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|Packet: LS Ack
2020-08-12:11:59:50.272332|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_PACKET|
(End of packet, OSPFv3 process Id = 1, VRF Name = default)
2020-08-12:11:59:45.072407|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|OSPF message details:
2020-08-12:11:59:45.072424|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|Header: Src IP:FE80::800:901:4C1:3751
Dst IP:FE80::800:901:40D:A274 Type: 2 Rtr Id:2.2.2.2 Area:0.0.0.1 Inst ID:0
2020-08-12:11:59:45.072443|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|Packet: DB Desc: If MTU:1500 Opt:19 I,M,MS
bits:0X01 Seq No:1020993
2020-08-12:11:59:45.072460|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_PACKET|
(End of packet, OSPFv3 process Id = 1, VRF Name = default)
2020-08-12:11:59:45.073077|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|OSPF message details:
2020-08-12:11:59:45.073096|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|Header: Src IP:FE80::800:901:4C1:3751
Dst IP:FE80::800:901:40D:A274 Type: 3 Rtr Id:1.1.1.1 Area:0.0.0.1 Inst ID:0
2020-08-12:11:59:45.073114|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|Packet: LS Req: First LS Type:8193 LS id:0.0.0.0
Adv Router:2.2.2.2
2020-08-12:11:59:45.073130|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_PACKET|
(End of packet, OSPFv3 process Id = 1, VRF Name = default)
2020-08-12:12:11:22.344623|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|OSPF message details:
2020-08-12:12:11:22.344635|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|Header: Src IP:FE80::800:901:4C1:3751
Dst IP:FE80::800:901:40D:A274 Type: 2 Rtr Id:2.2.2.2 Area:0.0.0.1 Inst ID:0
2020-08-12:12:11:22.344646|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_
PACKET|Packet: DB Desc: If MTU:1500 Opt:19 I,M,MS bits:0X07
Seq No:1728268
2020-08-12:12:11:22.344663|hpe-routing|LOG_DEBUG|AMM|-|OSPFV3-AF|OSPFv3_AF_PACKET|
(End of packet, OSPFv3 process Id = 1, VRF Name = default)
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ospfv3-af ipv4

`show ospfv3-af [<PROCESS-ID>] ipv4 [all-vrfs | vrf <VRF-NAME>]`

Description

Shows OSPFv3 AF information including area, state, and configuration information.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 AF process ID optionally to show OSPFv3 AF information for a particular OSPF process. Range: 1 to 65535.
`all-vrfs`	Select to show OSPFv3 AF information for all VRFs.
`vrf <VRF-NAME>`	Specifies the name of a VRF. The default VRF OSPFv3 AF process information is displayed if VRF is not mentioned. Default: default.

Example

Showing general OSPFv3 AF configurations:

```
switch# show ospfv3-af 200 ipv4
Address-family : IPv4 Unicast
------------------------------


VRF : default                             Process : 20200
-------------------------------------------------------------------------

Router ID              : 1.1.1.1          OSPFv3-AF              : Enabled
BFD                    : Disabled         SPF Start Interval     : 200 ms
SPF Hold Interval      : 1000 ms          SPF Max Wait Interval  : 5000 ms
LSA Start Interval     : 5000 ms          LSA Hold Interval      : 1000 ms
LSA Max Wait Interval  : 1000 ms          LSA Arrival Interval   : 1000 ms
External LSAs          : 0                Checksum Sum           : 0
ECMP                   : 4                Reference Bandwidth    : 100000 Mbps
Area Border            : Yes              AS Border              : No
GR Status              : Enabled          GR Interval            : 120 sec
GR State               : Inactive         GR Exit Status         : None
GR Helper              : Enabled          GR Strict LSA Check    : Enabled
GR Ignore Lost I/F     : Disabled         Internal Process ID    : 1
Summary address:
  prefix fd00::1/64, advertise, tag 10
```

```
Area      Total Active
----------------------
Normal      2      2
Stub        0      0

Area  : 0.0.0.0
----------------------------------
Area Type             : Normal       Status               : Active
Total Interfaces      : 1            Active Interfaces    : 1
Passive Interfaces    : 0            Loopback Interfaces  : 0
SPF calculation count : 4
Area ranges:
  fd00::1/64, inter-area, no-advertise
AH Authentication     : SHA1, SPI 256
Number of LSAs        : 5            Checksum Sum         : 99122

Area  : 0.0.0.1
----------------------------------
Area Type             : Normal       Status               : Active
Total Interfaces      : 1            Active Interfaces    : 1
Passive Interfaces    : 0            Loopback Interfaces  : 0

SPF Calculation Count : 4
Area ranges:
  fd00::1/64, inter-area, no-advertise
ESP Authentication    : SHA1
Encryption            : 3DES, SPI 256
Number of LSAs        : 5            Checksum Sum         : 99122
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ospfv3-af ipv4 border-routers

`show ospfv3-af [<PROCESS-ID>] ipv4 border-routers | [all-vrfs | vrf <VRF-NAME>]`

Description

Shows the OSPFv3 AF routing table entries for Area Border Router (ABR) and Autonomous System Border Router (ASBR).

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 AF process ID to show the OSPFv3 AF routing table entries for ABR and ASBR for the particular OSPFv3 AF process. Range: 1 to 65535.
`all-vrfs`	Select to show OSPFv3 AF border router information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Example

Showing OSPFv3 AF border routers information for VRF **vrf_red**:

```
switch# show ospfv3-af ipv4 border-routers vrf vrf_red

Address-family : IPv4 Unicast
------------------------------

VRF : vrf_red       Process ID : 1
             Internal Routing Table
------------------------------------------------------

Codes: i - Intra-area route, I - Inter-area route
   Router-ID Cost  Type    Area      SPF   Nexthop       Interface
i  1.1.1.1   1     ASBR    0.0.0.0   9     10.10.10.1    1/1/2
i  3.3.3.3   1     ASBR    1.1.1.1   9     20.20.20.1    ** tunnel1
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ospfv3-af ipv4 interface

`show ospfv3-af [`*<PROCESS-ID>*`] ipv4 interface [`*<INTERFACE-NAME>*`] | vlan` *<VLAN-ID>* `| lag` *<LAG-ID>* `| loopback` *<LOOPBACK-ID>* `| tunnel` *<TUNNEL-ID>*`]`

```
[brief] [all-vrfs | vrf <VRF-NAME>]
```

Description

Shows information about OSPFv3 AF enabled IPv4 interfaces.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 AF process ID optionally to show the OSPFv3 AF enabled interfaces for the particular OSPFv3 AF process. Range: 1 to 65535.
<INTERFACE-NAME>	Selects to show information only for the specified OSPFv3 AF enabled interface.
`vlan` *<VLAN-ID>*	Specifies the vlan interface name.
`lag` *<LAG-ID>*	Specifies the lag interface name.
`loopback` *<LOOPBACK-ID>*	Specifies the loopback interface name.
`tunnel` *<TUNNEL-ID>*	Specifies the tunnel interface name.
`brief`	Include this parameter to display a brief overview of the following OSPFv3 AF configuration information. ■ Interface: OSPFv3 AF interface name. ■ Area: OSPFv3 AF area ID. ■ Cost: The metric OSPFv3 AF uses to judge a path's feasibility, calculated as (reference bandwidth / interface bandwidth). ■ State: Indicates if the interface is a designated router (**Dr**) or a backup designated router (**Backup-dr**). ■ Status: Indicates if the interface is **up** or **down**. ■ Flags: P - Passive A - Active.
`all-vrfs`	Select to show OSPFv3 AF-enabled interface information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Examples

Showing OSPFv3 AF information for all interfaces in default VRF:

```
switch# show ospfv3-af ipv4 interface
Codes: DR - Designated router BDR - Backup Designated router

Address-family : IPv4 Unicast
------------------------------

Interface 1/1/1 is up, Line Protocol is up, Instance Id 64
------------------------------------------------------------
 VRF                 : default                          Process             : 1
 IP Address          : 1.1.1.2                          Area                :
0.0.0.0
 Status              : Up                               Network Type        :
Broadcast
 Hello Interval      : 10    sec                        Dead Interval       : 40
   sec
```

```
 Transit Delay        : 1     sec                      Retransmit Interval  : 5
   sec
 BFD                  : Disabled                       Link Speed           :
1000 Mbps
 Cost Configured      : 1                              Cost Calculated      : 4
 State/Type           : BDR                            Router Priority      : 1
 DR                   : 3.3.3.3                        BDR                  :
1.1.1.3
 Link LSAs            : 2                              Checksum Sum         :
26868
 AH Authentication    : SHA1, SPI 256 (Area)           Passive              : Yes

Interface 1/1/2 is up, Line Protocol is up, Instance Id 64
------------------------------------------------------------
 VRF                  : default                        Process              : 1
 IP Address           : 10.10.10.2                     Area                 :
0.0.0.0
 Status               : Up                             Network Type         :
Broadcast
 Hello Interval       : 10    sec                      Dead Interval        : 40
   sec
 Transit Delay        : 1     sec                      Retransmit Interval  : 5
   sec
 Link Speed           : 1000 Mbps
 Cost Configured      : 1                              Cost Calculated      : 4
 State/Type           : BDR                            Router Priority      : 1
 DR                   : 10.10.10.2                     BDR                  :
10.10.10.1
 Link LSAs            : 0                              Checksum Sum         : 0
 AH Authentication    : SHA1, SPI 256 (Area)           Passive              : No

Interface tunnel1 is up, Line Protocol is up, Instance Id 64
------------------------------------------------------------
 VRF                  : default                        Process              : 1
 IPv6 Address         : 10.10.10.1                     Area                 :
0.0.0.1
 Status               : Up                             Network Type         :
Broadcast
 Hello Interval       : 10    sec                      Dead Interval        : 40
   sec
 Transit Delay        : 1     sec                      Retransmit Interval  : 5
   sec
 BFD                  : Disabled                       Link Speed           :
1000 Mbps
 Cost Configured      : 1                              Cost Calculated      : 4
 State/Type           : DR                             Router Priority      : 1
 DR                   : 10.10.10.2                     BDR                  :
10.10.10.1
 Link LSAs            : 1                              Checksum Sum         :
9168
 AH Authentication    : SHA1, SPI 256 (Area)           Passive              : Yes

Interface vlan10 is up, Line Protocol is up, Instance Id 64
------------------------------------------------------------
 VRF                  : default                        Process              : 1
 IPv6 Address         : 10.10.10.3                     Area                 :
0.0.0.1
 Status               : Up                             Network Type         :
Broadcast
 Hello Interval       : 10    sec                      Dead Interval        : 40
   sec
 Transit Delay        : 1     sec                      Retransmit Interval  : 5
```

```
   sec
 BFD                  : Disabled                          Link Speed           :
1000 Mbps
 Cost Configured      : 1                                 Cost Calculated      : 4
 State/Type           : DR                                Router Priority      : 1
 DR                   : 10.10.10.2                        BDR                  :
10.10.10.1
 Link LSAs            : 1                                 Checksum Sum         :
9168
 AH Authentication    : MD5, SPI 257                      Passive              : No

Interface vlan20 is up, Line Protocol is up, Instance Id 64
------------------------------------------------------------

 VRF                  : default                           Process              : 1
 IPv6 Address         : 1.1.1.1                           Area                 :
0.0.0.1
 Status               : Up                                Network Type         :
Broadcast
 Hello Interval       : 10    sec                         Dead Interval        : 40
   sec
 Transit Delay        : 1     sec                         Retransmit Interval  : 5
   sec
 Link Speed           : 1000 Mbps
 Cost Configured      : 1                                 Cost Calculated      : 4
 State/Type           : DR                                Router Priority      : 1
 DR                   : 1.1.1.3                           BDR                  : NA
 Link LSAs            : 1                                 Checksum Sum         :
9168
 AH Authentication    : MD5, SPI 257                      Passive              : No
```

Showing overview information for OSPFv3 AF enabled interfaces for all VRFs in brief:

```
switch# show ospfv3-af ipv4 interface brief all-vrfs
Address-family : IPv4 Unicast
------------------------------

OSPFv3 AF Process ID 1 VRF default
==================================

Total Number of Interfaces: 2

Flags: P - Passive  A - Active

Interface     Area          Cost      State            Status    Flags
--------------------------------------------------------------------------
1/1/1         0.0.0.0       1         BDR              Up        P
vlan10        0.0.0.1       1         DR               Up        A

OSPFv3 AF Process ID 2 VRF red
==============================

Total Number of Interfaces: 1

Flags: P - Passive  A - Active

Interface     Area          Cost      State            Status    Flags
--------------------------------------------------------------------------
1/1/2         0.0.0.1       1         BDR              Up        P
tunnel1       0.0.0.0       100       DR               Up        A
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ospfv3-af ipv4 lsdb

```
show ospfv3-af [<PROCESS-ID>] ipv4 lsdb [all-vrfs | vrf <VRF-NAME>] [router | network |
inter-area-prefix | inter-area-router | as-external | nssa-external | link | intra-area-
prefix]
   [area <AREA-ID>] [lsid <LINK-STATE-ID>] [adv-router <ROUTER-ID> | self]
show ospfv3-af [<PROCESS-ID>] ipv4 lsdb [all-vrfs | vrf <VRF-NAME>] database-summary
```

Description

Shows the OSPFv3 AF link state database summary for different OSPFv3 AF LSAs (Link State Advertisement). Use the parameters to get information for a particular LSA.

Parameter	Description
<PROCESS-ID>	Enter an OSPFv3 AF process ID to display general OSPF information for a particular OSPFv3 AF process. Range: 1 to 65535.
all-vrfs|vrf <vrf-name>	Select **all-vrfs** to display general OSPF information for all VRFs, or use the **vrf <*VRF-NAME*>** option to display information for a specific VRF. Optionally select one of the following parameters to filter the link state database information.
as-external	Show external link states (LSA type 5)
database-summary	Select to display the count of each type of LSA and each area in the database. **NOTE:** The **database-summary** parameter does not support the **area <area-id>**, **lsid <link-state-id>** or **adv-router {<router-id>|self}** parameters.
inter-area-prefix	Show inter-area prefix link states (LSA type 3)

Parameter	Description
`inter-area-router`	Show inter-area router link states (LSA type 4)
`intra-area-prefix`	Show intra-area prefix link states (LSA type 9
`link`	Show link states (LSA type 8)
`network`	Show network LSAs (LSA type 2).
`nssa-external`	Show NSSA external link states (LSA type 7).
`router`	Show router LSAs (LSA type 1).
`area <AREA-ID>`	Select to display information filtered for the specified area in one of the following formats. OSPFv3 AF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPFv3 AF area identifier in decimal format. Value: 0 to 4294967295.
`lsid <LINK-STATE-ID>`	Select to display information filtered by link state identifier specified in IPv4 address format (A.B.C.D).
`adv-router {<ROUTER-ID>|self}`	Select to display link states for a particular advertising router. Specify either a Router ID of the advertising router or specify **self** to show self-originated link states.

Examples

Showing OSPFv3 AF link state database (LSDB) general information:

```
switch# show ospfv3-af ipv4 lsdb
Address-family : IPv4 Unicast
------------------------------

OSPF Router with ID (50.50.50.50) (Process ID 1 VRF default)
============================================================

Router Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID          Link Count  Bits
-------------------------------------------------------------------------------
40.40.40.40     930     0x80000004  0x2ea1      0             3           None
50.50.50.50     935     0x80000002  0x8b52      0             1           E
60.60.60.60     943     0x800003c5  0x9854      0             2           None

Network Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Router Count
-------------------------------------------------------------------------
60.60.60.60     944     0x80000001  0x7179      1360007168  2
50.50.50.50     935     0x80000001  0x516a      19          1

Inter Area Prefix Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
-----------------------------------------------------------------
40.40.40.40     929     0x80000001  0x2498      131072      11.0.0.0/8
50.50.50.50     928     0x80000001  0x5b2f      65536       100.100.0.0/16
```

```
Inter Area Router Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Destination Router ID
-------------------------------------------------------------------------------
40.40.40.40     929     0x80000001  0x2498      1           33.33.33.33

AS External Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
-------------------------------------------------------------------
40.40.40.40     264     0x80000001  0x24cc4     1           10.10.0.0/16
40.40.40.40     675     0x80000001  0x5b00f     2           11.11.0.0/16

NSSA External Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
-------------------------------------------------------------------
3.3.3.3         264     0x80000001  0x24ac2     1           200.200.0.0/16

Link-local Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Interface
-------------------------------------------------------------------
50.50.50.50     264     0x80000001  0x653c4     19          1/1/1

Intra Area Prefix Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------------
-------------
ADV Router      Age     Seq#        Checksum    LSID        Referenced LS Type
Referenced LSID
-------------------------------------------------------------------------------
-------------
50.50.50.50     263     0x80000001  0x1da34     1           0x2001              0
50.50.50.50     264     0x80000001  0x2a45d     1           0x2002              19
```

Showing AS external link states:

```
switch# show ospfv3-af ipv4 lsdb as-external
Address-family : IPv4 Unicast
-----------------------------

OSPF Router with ID (60.60.60.60) (Process ID 1 VRF default)
============================================================

AS External Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
-------------------------------------------------------------------
40.40.40.40     264     0x80000001  0x24cc4     1           10.10.0.0/16
40.40.40.40     675     0x80000001  0x5b00f     2           11.11.0.0/16
```

Showing the LSDB database summary:

```
switch# show ospfv3-af ipv4 lsdb database-summary
Address-family : IPv4 Unicast
-----------------------------

OSPF Router with ID (10.1.1.1) (Process ID 1 VRF default)
```

```
============================================================

Area 0.0.0.0 database summary
---------------------------
LSA Type             Count
---------------------------
Router               2
Network              1
Inter Area Prefix    1
Inter Area Router    0
NSSA External        0
Link                 3
Intra Area Prefix    3
---------------------------
Total                10

Process 1 database summary
---------------------------
LSA Type             Count
---------------------------
Router               2
Network              1
Inter Area Prefix    1
Inter Area Router    0
AS External          2
NSSA External        0
Link                 3
Intra Area Prefix    3
---------------------------
Total                12
```

Showing inter-area prefix LSAs:

```
switch# show ospfv3-af ipv4 lsdb inter-area-prefix
Address-family : IPv4 Unicast
------------------------------

OSPF Router with ID (6.6.6.6) (Process ID 1 VRF default)
============================================================

Inter Area Prefix Link State Advertisements (Area 0.0.0.0)
------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
------------------------------------------------------------------
40.40.40.40     929     0x80000001  0x2498      131072      1.0.0.0/8
50.50.50.50     928     0x80000001  0x5b2f      65536       11.11.0.0./16
```

Showing inter-area router LSAs:

```
switch# show ospfv3-af ipv4 lsdb inter-area-router
Address-family : IPv4 Unicast
------------------------------

OSPF Router with ID (60.60.60.60) (Process ID 1 VRF default)
============================================================

Inter Area Prefix Link State Advertisements (Area 0.0.0.0)
------------------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Destination Router ID
```

```
-------------------------------------------------------------------------------
40.40.40.40     929     0x80000001  0x2498      1           33.33.33.33
```

Showing intra-area prefix LSAs:

```
switch# show ospfv3-af ipv4 lsdb intra-area-prefix
Address-family : IPv4 Unicast
-----------------------------

OSPF Router with ID (0.0.0.0) (Process ID 1 VRF default)
========================================================

Intra Area Prefix Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------------
-----------------
ADV Router      Age     Seq#        Checksum    LSID        Referenced LS Type
 Referenced LSID
-------------------------------------------------------------------------------
-----------------
50.50.50.50     263     0x80000001  0x1da34      1          0x2001
 0
50.50.50.50     264     0x80000001  0x2a45d      1          0x2002
 19
```

Showing link LSAs:

```
switch# show ospfv3-af ipv4 lsdb link
Address-family : IPv4 Unicast
-----------------------------

OSPF Router with ID (2.2.2.1) (Process ID 1 VRF default)
========================================================

Link-local Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Interface
-------------------------------------------------------------------
50.50.50.50     264     0x80000001  0x653c4     19          1/1/1

Link-local Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Interface
-------------------------------------------------------------------
50.50.50.50     264     0x80000001  0xad324     20          1/1/2
```

Showing network LSAs:

```
switch# show ospfv3-af ipv4 lsdb network
Address-family : IPv4 Unicast
-----------------------------

OSPF Router with ID (50.50.50.50) (Process ID 1 VRF default)
============================================================

Network Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Router Count
```

```
-------------------------------------------------------------------------
60.60.60.60     944     0x80000001  0x7179      1360007168  2
50.50.50.50     935     0x80000001  0x516a      19          1
```

Showing NSSA external link states:

```
switch# show ospfv3-af ipv4 lsdb nssa-external
Address-family : IPv4 Unicast
------------------------------

OSPF Router with ID (2.2.2.1) (Process ID 1 VRF default)
=========================================================

NSSA External Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
-------------------------------------------------------------------
3.3.3.3         264     0x80000001  0x24ac2     1           200.200.0.0/16
```

Showing router LSAs:

```
switch# show ospfv3-af ipv4 lsdb router
Address-family : IPv4 Unicast
------------------------------

OSPF Router with ID (50.50.50.50) (Process ID 1 VRF default)
=========================================================

Router Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Link Count    Bits
-------------------------------------------------------------------------------
40.40.40.40     930     0x80000004  0x2ea1      0           3             None
50.50.50.50     935     0x80000002  0x8b52      0           1             E
60.60.60.60     943     0x800003c5  0x9854      0           2             None
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300	Operator (>) or Manager (#)	

Platforms	Command context	Authority
9300S 10000		

show ospfv3-af ipv4 neighbors

```
show ospfv3-af [<PROCESS-ID>] ipv4 neighbors [<NEIGHBOR-ID>]
   [interface <INTERFACE-NAME> | vlan <VLAN-ID> | lag <LAG-ID> | loopback <LOOPBACK-ID> |
   tunnel <TUNNEL-ID>]
   [detail | summary] [all-vrfs | vrf <VRF-NAME>]
```

Description

Shows information about OSPFv3 AF neighbors.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 AF process ID to show OSPFv3 AF neighbor information for the particular OSPFv3 AF process. Range: 1 to 65535.
`neighbors` *<NEIGHBOR-ID>*	Shows information about a particular neighbor, specified in IPv4 format (A.B.C.D).
`interface` *<INTERFACE-NAME>*	Shows neighbor information only for the specified interface.
`vlan` *<VLAN-ID>*	Specifies the vlan interface name.
`lag` *<LAG-ID>*	Specifies the lag interface name.
`loopback` *<LOOPBACK-ID>*	Specifies the loopback interface name.
`tunnel` *<TUNNEL-ID>*	Specifies the tunnel interface name.
`detail`	Shows detailed information for the neighbors.
`summary`	Shows summary information for the neighbors.
`all-vrfs`	Shows neighbor information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Examples

Showing OSPFv3 AF neighbor information for the specified interface:

```
switch# show ospfv3-af ipv4 neighbors 3.3.3.3
Address-family : IPv4 Unicast
------------------------------

VRF : default                          Process : 1
-------------------------------------------------------------
Router-Id          : 3.3.3.3           Area              : 0.0.0.0
Interface          : 1/1/1             Address           : 1.1.1.1
State              : FULL              Neighbor Priority : 1
Dead Timer Due     : 00:00:36          Options           : 0x13
Time since last state change : 00h:14m:45s
```

Showing detail OSPFv3 AF neighbor information:

```
switch# show ospfv3-af ipv4 neighbors 3.3.3.3 detail
Address-family : IPv4 Unicast
-----------------------------
VRF : default                           Process : 1
-------------------------------------------------------------
Router-Id         : 3.3.3.3             Area              : 0.0.0.0
Interface         : 1/1/1               Address           : 1.1.1.1
State             : FULL                Neighbor Priority : 1
DR                : 3.3.3.3             BDR               : 1.1.1.3
Dead Timer Due    : 00:00:36            Options           : 0x13
Retransmission Queue Length  : 0
Time Since Last State Change : 00h:14m:45s
```

Showing OSPFv3 AF neighbor information for interface **1/1/1**:

```
switch# show ospfv3-af ipv4 neighbors 3.3.3.3 interface 1/1/1
Address-family : IPv4 Unicast
-----------------------------
VRF : default                           Process : 1
-------------------------------------------------------------
Router-Id         : 3.3.3.3             Area              : 0.0.0.0
Interface         : 1/1/1               Address           : 1.1.1.1
State             : FULL                Neighbor Priority : 1
Dead Timer Due    : 00:00:36            Options           : 0x13
Time Since Last State Change : 00h:14m:45s
```

Showing OSPFv3 AF neighbor information for tunnel interface:

```
switch# show ospfv3-af ipv4 neighbors 3.3.3.3
Address-family : IPv4 Unicast
-----------------------------

VRF : default                           Process : 1
-------------------------------------------------------------
Router-Id         : 3.3.3.3             Area              : 0.0.0.0
Interface         : tunnel1             Address           : 1.1.1.1
State             : FULL                Neighbor Priority : 1
Dead Timer Due    : 00:00:36            Options           : 0x13
Time since last state change : 00h:14m:45s
```

Showing detail OSPFv3 AF neighbor information for tunnel interface:

```
switch# show ospfv3-af ipv4 neighbors 3.3.3.3 detail
Address-family : IPv4 Unicast
-----------------------------

VRF : default                           Process : 1
-------------------------------------------------------------
Router-Id         : 3.3.3.3             Area              : 0.0.0.0
Interface         : tunnel1             Address           : 1.1.1.1
State             : FULL                Neighbor Priority : 1
DR                : 3.3.3.3             BDR               : 1.1.1.3
Dead Timer Due    : 00:00:36            Options           : 0x13
Retransmission Queue Length  : 0
Time Since Last State Change : 00h:14m:45s
```

Showing summary OSPFv3 AF neighbors information for a specific neighbor for all VRFs:

```
switch# show ospfv3-af ipv4 neighbors 3.3.3.3 summary all-vrfs
Address-family : IPv4 Unicast
------------------------------

OSPFv3 AF Process ID 1 VRF default, Neighbor Summary
================================================

Interface    Down Attempt Init TwoWay ExStart Exchange Loading Full Total
-------------------------------------------------------------------------
1/1/1        0    0       0    0      0       0        0       1    1
Total        0    0       0    0      0       0        0       1    1

OSPFv3 AF Process ID 1 VRF red, Neighbor Summary
============================================

Interface    Down Attempt Init TwoWay ExStart Exchange Loading Full Total
-------------------------------------------------------------------------
1/1/2        0    0       0    0      0       0        0       1    1
Total        0    0       0    0      0       0        0       1    1
```

Showing OSPFv3 AF neighbors information for VRF red:

```
switch# show ospfv3-af ipv4 neighbors vrf red
Address-family : IPv4 Unicast
------------------------------

OSPFv3 AF Process ID 2 VRF red
==============================

Total Number of Neighbors: 1

Neighbor ID      Priority  State             Interface
-----------------------------------------------------
4.4.4.4          1         FULL/DR           1/1/2
  Neighbor address 1.1.1.1
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300	Operator (>) or Manager (#)	

Platforms	Command context	Authority
9300S 10000		

show ospfv3-af ipv4 routes

```
show ospfv3-af [<PROCESS-ID>] ipv4 routes [<PREFIX/LENGTH>]
   [all-vrfs | vrf <VRF-NAME>]
```

Description

Shows the OSPFv3 AF routing table information.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 AF process ID that shows information from the OSPFv3 AF routing table for the particular OSPFv3 AF process. Range: 1 to 65535.
<PREFIX/LENGTH>	Specifies an IP address in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.
`all-vrfs`	Select to show OSPFv3 AF routing table information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Examples

Showing OSPFv3 AF routing table information:

```
switch# show ospfv3-af ipv4 routes
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

Address-family : IPv4 Unicast
------------------------------

OSPFv3 AF Process ID 1 VRF default, Routing Table
-------------------------------------------------

Total Number of OSPFv3 AF Routes : 3

11.11.0.0/16       (i) area:0.0.0.0
     directly attached to interface 1/1/1, cost 1 distance 110
200.200.0.0/16     (i) area:0.0.0.0
     directly attached to interface tunnel1, cost 100 distance 110
1.1.1.1/132        (i) area:0.0.0.1
     directly attached to interface vlan10, cost 1 distance 110
```

Showing OSPFv3 AF routing table information for VRF red:

```
switch# show ospfv3-af ipv4 routes vrf red

Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2
```

```
Address-family : IPv4 Unicast
------------------------------

OSPFv3 AF Process ID 2 VRF red, Routing Table
---------------------------------------------

Total Number of OSPFv3 AF Routes : 1

2.2.2.1/24           (i) area:0.0.0.1
     directly attached to interface 1/1/2, cost 1 distance 110
```

Showing OSPFv3 AF routing table information for destination prefix 2.2.2.1/24:

```
switch# show ospfv3-af 1 ipv4 routes 2.2.2.1/24
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

Address-family : IPv4 Unicast
------------------------------

OSPFv3 AF Process ID 1 VRF default, Routing Table for prefixes 2.2.2.1/24
--------------------------------------------------------------------------

Total Number of OSPFv3 AF Routes : 1

2.2.2.1/24          (i) area:0.0.0.1
     directly attached to interface vlan10, cost 1 distance 110
```

Showing OSPFv3 AF routing table information for destination prefix 200.0.0.0/8:

```
switch# show ospfv3-af 1 ipv4 routes 200.0.0.0/8
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

Address-family : IPv4 Unicast
------------------------------

OSPFv3 AF Process ID 1 VRF default, Routing Table for prefixes 200.0.0.0/8
--------------------------------------------------------------------------

Total Number of OSPFv3 AF Routes : 1

200.0.0.0/8          (i) area:0.0.0.0
     directly attached to interface tunnel1, cost 100 distance 110
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A,	Operator (>) or	

Platforms	Command context	Authority
S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Manager (#)	

show ospfv3-af ipv4 statistics interface

```
show ospfv3-af [<PROCESS-ID>] ipv4 statistics interface [<INTERFACE-NAME>] | vlan <VLAN-
ID> | lag <LAG-ID> | loopback <LOOPBACK-ID> | tunnel <TUNNEL-ID>]
   [brief] [all-vrfs | vrf <VRF-NAME>]
```

Description

Shows information about OSPFv3 AF enabled IPv4 interfaces.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 AF process ID optionally to show the OSPFv3 AF enabled interfaces for the particular OSPFv3 AF process. Range: 1 to 65535.
<INTERFACE-NAME>	Selects to show information only for the specified OSPFv3 AF enabled interface.
`vlan` *<VLAN-ID>*	Specifies the vlan interface name.
`lag` *<LAG-ID>*	Specifies the lag interface name.
`loopback` *<LOOPBACK-ID>*	Specifies the loopback interface name.
`tunnel` *<TUNNEL-ID>*	Specifies the tunnel interface name.
`brief`	Include this parameter to display a brief overview of the following OSPFv3 AF configuration information. ■ Interface: OSPFv3 AF interface name. ■ Area: OSPFv3 AF area ID. ■ Cost: The metric OSPFv3 AF uses to judge a path's feasibility, calculated as (reference bandwidth / interface bandwidth). ■ State: Indicates if the interface is a designated router (**Dr**) or a backup designated router (**Backup-dr**). ■ Status: Indicates if the interface is **up** or **down**. ■ Flags: P - Passive A - Active.
`all-vrfs`	Select to show OSPFv3 AF-enabled interface information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Examples

Showing OSPFv3 AF information for all interfaces in default VRF:

```
switch# show ospfv3-af ipv4 interface
Codes: DR - Designated router BDR - Backup Designated router

Address-family : IPv4 Unicast
------------------------------

Interface 1/1/1 is up, Line Protocol is up, Instance Id 64
------------------------------------------------------------
 VRF                 : default                       Process             : 1
 IP Address          : 1.1.1.2                       Area                :
0.0.0.0
 Status              : Up                            Network Type        :
Broadcast
 Hello Interval      : 10    sec                     Dead Interval       : 40
   sec
 Transit Delay       : 1     sec                     Retransmit Interval : 5
   sec
 BFD                 : Disabled                      Link Speed          :
1000 Mbps
 Cost Configured     : 1                             Cost Calculated     : 4
 State/Type          : BDR                           Router Priority     : 1
 DR                  : 3.3.3.3                       BDR                 :
1.1.1.3
 Link LSAs           : 2                             Checksum Sum        :
26868
 AH Authentication   : SHA1, SPI 256 (Area)          Passive             : Yes

Interface 1/1/2 is up, Line Protocol is up, Instance Id 64
------------------------------------------------------------
 VRF                 : default                       Process             : 1
 IP Address          : 10.10.10.2                    Area                :
0.0.0.0
 Status              : Up                            Network Type        :
Broadcast
 Hello Interval      : 10    sec                     Dead Interval       : 40
   sec
 Transit Delay       : 1     sec                     Retransmit Interval : 5
   sec
 Link Speed          : 1000 Mbps
 Cost Configured     : 1                             Cost Calculated     : 4
 State/Type          : BDR                           Router Priority     : 1
 DR                  : 10.10.10.2                    BDR                 :
10.10.10.1
 Link LSAs           : 0                             Checksum Sum        : 0
 AH Authentication   : SHA1, SPI 256 (Area)          Passive             : No

Interface tunnel1 is up, Line Protocol is up, Instance Id 64
------------------------------------------------------------
 VRF                 : default                       Process             : 1
 IPv6 Address        : 10.10.10.1                    Area                :
0.0.0.1
 Status              : Up                            Network Type        :
Broadcast
 Hello Interval      : 10    sec                     Dead Interval       : 40
   sec
 Transit Delay       : 1     sec                     Retransmit Interval : 5
   sec
 BFD                 : Disabled                      Link Speed          :
1000 Mbps
 Cost Configured     : 1                             Cost Calculated     : 4
 State/Type          : DR                            Router Priority     : 1
```

```
 DR                   : 10.10.10.2                      BDR                  :
10.10.10.1
 Link LSAs            : 1                               Checksum Sum         :
9168
 AH Authentication    : SHA1, SPI 256 (Area)            Passive              : Yes

Interface vlan10 is up, Line Protocol is up, Instance Id 64
------------------------------------------------------------
 VRF                  : default                         Process              : 1
 IPv6 Address         : 10.10.10.3                      Area                 :
0.0.0.1
 Status               : Up                              Network Type         :
Broadcast
 Hello Interval       : 10    sec                       Dead Interval        : 40
   sec
 Transit Delay        : 1     sec                       Retransmit Interval  : 5
   sec
 BFD                  : Disabled                        Link Speed           :
1000 Mbps
 Cost Configured      : 1                               Cost Calculated      : 4
 State/Type           : DR                              Router Priority      : 1
 DR                   : 10.10.10.2                      BDR                  :
10.10.10.1
 Link LSAs            : 1                               Checksum Sum         :
9168
 AH Authentication    : MD5, SPI 257                    Passive              : No

Interface vlan20 is up, Line Protocol is up, Instance Id 64
------------------------------------------------------------

 VRF                  : default                         Process              : 1
 IPv6 Address         : 1.1.1.1                         Area                 :
0.0.0.1
 Status               : Up                              Network Type         :
Broadcast
 Hello Interval       : 10    sec                       Dead Interval        : 40
   sec
 Transit Delay        : 1     sec                       Retransmit Interval  : 5
   sec
 Link Speed           : 1000 Mbps
 Cost Configured      : 1                               Cost Calculated      : 4
 State/Type           : DR                              Router Priority      : 1
 DR                   : 1.1.1.3                         BDR                  : NA
 Link LSAs            : 1                               Checksum Sum         :
9168
 AH Authentication    : MD5, SPI 257                    Passive              : No
```

Showing overview information for OSPFv3 AF enabled interfaces for all VRFs in brief:

```
switch# show ospfv3-af ipv4 interface
Address-family : IPv4 Unicast
------------------------------

OSPFv3 AF Process ID 1 VRF default
==================================

Total Number of Interfaces: 2

Flags: P - Passive  A - Active

Interface     Area          Cost      State            Status    Flags
```

```
--------------------------------------------------------------------------
1/1/1          0.0.0.0       1          BDR              Up         P
vlan10         0.0.0.1       1          DR               Up         A

OSPFv3 AF Process ID 2 VRF red
==============================

Total Number of Interfaces: 1

Flags: P - Passive  A - Active

Interface      Area          Cost       State            Status     Flags
--------------------------------------------------------------------------
1/1/2          0.0.0.1       1          BDR              Up         P
tunnel1        0.0.0.0       100        DR               Up         A
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ospfv3-af ipv6

`show ospfv3-af [<PROCESS-ID>] ipv6 [all-vrfs | vrf <VRF-NAME>]`

Description

Shows OSPFv3 AF information including area, state, and configuration information.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 AF process ID optionally to show OSPFv3 AF information for a particular OSPF process. Range: 1 to 65535.
`all-vrfs`	Select to show OSPFv3 AF information for all VRFs.
`vrf <VRF-NAME>`	Specifies the name of a VRF. The default VRF OSPFv3 AF process information is displayed if VRF is not mentioned. Default: default.

Example

Showing general OSPFv3 AF configurations:

```
switch# show ospfv3-af 200 ipv6
Address-family : IPv6 Unicast
------------------------------

VRF : default                              Process : 20200
-------------------------------------------------------------------------

Router ID              : 1.1.1.1           OSPFv3-AF              : Enabled
BFD                    : Disabled          SPF Start Interval     : 200 ms
SPF Hold Interval      : 1000 ms           SPF Max Wait Interval  : 5000 ms
LSA Start Interval     : 5000 ms           LSA Hold Interval      : 1000 ms
LSA Max Wait Interval  : 1000 ms           LSA Arrival Interval   : 1000 ms
External LSAs          : 0                 Checksum Sum           : 0
ECMP                   : 4                 Reference Bandwidth    : 100000 Mbps
Area Border            : Yes               AS Border              : No
GR Status              : Enabled           GR Interval            : 120 sec
GR State               : Inactive          GR Exit Status         : None
GR Helper              : Enabled           GR Strict LSA Check    : Enabled
GR Ignore Lost I/F     : Disabled          Internal Process ID    : 1
Summary address:
  prefix fd00::1/64, advertise, tag 10

Area      Total Active
----------------------
Normal      2      2
Stub        0      0

Area  : 0.0.0.0
---------------------------------
Area Type              : Normal       Status               : Active
Total Interfaces       : 1            Active Interfaces    : 1
Passive Interfaces     : 0            Loopback Interfaces  : 0
SPF calculation count : 4
Area ranges:
  fd00::1/64, inter-area, no-advertise
AH Authentication      : SHA1, SPI 256
Number of LSAs         : 5            Checksum Sum         : 99122

Area  : 0.0.0.1
---------------------------------
Area Type              : Normal       Status               : Active
Total Interfaces       : 1            Active Interfaces    : 1
Passive Interfaces     : 0            Loopback Interfaces  : 0

SPF Calculation Count : 4
Area ranges:
  fd00::1/64, inter-area, no-advertise
ESP Authentication     : SHA1
Encryption             : 3DES, SPI 256
Number of LSAs         : 5            Checksum Sum         : 99122
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ospfv3-af ipv6 border-routers

```
show ospfv3-af [<PROCESS-ID>] ipv6 border-routers | [all-vrfs | vrf <VRF-NAME>]
```

Description

Shows the OSPFv3 AF routing table entries for Area Border Router (ABR) and Autonomous System Border Router (ASBR).

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 AF process ID to show the OSPFv3 AF routing table entries for ABR and ASBR for the particular OSPFv3 AF process. Range: 1 to 65535.
`all-vrfs`	Select to show OSPFv3 AF border router information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Example

Showing OSPFv3 AF border routers information for VRF **vrf_red**:

```
switch# show ospfv3-af ipv6 border-routers vrf vrf_red

Address-family : IPv6 Unicast
------------------------------

VRF : vrf_red       Process ID : 1
             Internal Routing Table
------------------------------------------------------

Codes: i - Intra-area route, I - Inter-area route
   Router-ID Cost  Type    Area    SPF  Nexthop                      Interface
i  1.1.1.1   1     ASBR    0.0.0.0 9    fe80::7272:cfff:fe9a:a15d   1/1/2
i  3.3.3.3   1     ASBR    1.1.1.1 9    fe80::7272:cfff:fe1f:d80   ** tunnel1
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ospfv3-af ipv6 interface

```
show ospfv3-af [<PROCESS-ID>] ipv6 interface [<INTERFACE-NAME>] | vlan <vlan id> |
lag <lag id> | loopback <loopback id> | tunnel <tunnel id>] [brief]
[all-vrfs | vrf <VRF-NAME>]
```

Description

Shows information about OSPFv3 AF enabled interfaces.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 AF process ID optionally to show the OSPFv3 AF enabled interfaces for the particular OSPFv3 AF process. Range: 1 to 65535.
<INTERFACE-NAME>	Selects to show information only for the specified OSPFv3 AF enabled interface.
`vlan` *<VLAN-ID>*	Specifies the vlan interface name.
`lag` *<LAG-ID>*	Specifies the lag interface name.
`loopback` *<LOOPBACK-ID>*	Specifies the loopback interface name.
`tunnel` *<TUNNEL-ID>*	Specifies the tunnel interface name.
`brief`	Include this parameter to display a brief overview of the following OSPFv3 AF configuration information. ■ Interface: OSPFv3 AF interface name. ■ Area: OSPFv3 AF area ID. ■ Cost: The metric OSPFv3 AF uses to judge a path's feasibility, calculated as (reference bandwidth / interface bandwidth). ■ State: Indicates if the interface is a designated router (**Dr**) or a backup designated router (**Backup-dr**). ■ Status: Indicates if the interface is **up** or **down**. ■ Flags: P - Passive A - Active.
`all-vrfs`	Select to show OSPFv3 AF-enabled interface information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Examples

Showing OSPFv3 AF information for all interfaces in default VRF:

```
switch# show ospfv3-af ipv6 interface
Codes: DR - Designated router  BDR - Backup Designated router

Interface 1/1/1 is Up, Line Protocol is Up, Instance ID 10
-------------------------------------------
VRF                  : default                        Process              : 1
IPv6 address         : fe80::7272:cfff:fe70:67ae      Area                 :
0.0.0.0
Status               : Up                             Network Type         :
Broadcast
Hello Interval       : 10    sec                      Dead Interval        : 40
 sec
Transit Delay        : 1     sec                      Retransmit Interval  : 5
 sec
BFD                  : Disabled                       Link Speed           : 1000
Mbps
Cost Configured      : 1                              Cost Calculated      : 4
State/Type           : BDR                            Router Priority      : 1
DR                   : 3.3.3.3                        BDR                  :
1.1.1.3
Link LSAs            : 2                              Checksum Sum         :
26868
AH Authentication    : SHA1, SPI 256 (Area)           Passive              : Yes

Codes: DR - Designated router  BDR - Backup Designated router

Interface 1/1/2 is Up, Line Protocol is Up, Instance ID 70
-------------------------------------------
VRF                  : default                        Process              : 1
IPv6 address         : fe80::7272:cfff:fe70:67ae      Area                 :
0.0.0.0
Status               : Up                             Network Type         :
Broadcast
Hello Interval       : 10    sec                      Dead Interval        : 40
 sec
Transit Delay        : 1     sec                      Retransmit Interval  : 5
 sec
BFD                  : Disabled                       Link Speed           : 1000
Mbps
Cost Configured      : 1                              Cost Calculated      : 4
State/Type           : BDR                            Router Priority      : 1
DR                   : 10.10.10.2                     BDR                  :
10.10.10.1
Link LSAs            : 0                              Checksum Sum         : 0
AH Authentication    : SHA1, SPI 256 (Area)           Passive              : No

Codes: DR - Designated router  BDR - Backup Designated router

Interface tunnel1 is Up, Line Protocol is Up, Instance ID 70
-------------------------------------------
VRF                  : default                        Process              : 1
IPv6 address         : fe80::7272:cfff:fe70:679ae     Area                 :
0.0.0.1
Status               : Up                             Network Type         :
Broadcast
Hello Interval       : 10    sec                      Dead Interval        : 40
 sec
Transit Delay        : 1     sec                      Retransmit Interval  : 5
 sec
```

```
BFD                  : Disabled                          Link Speed           : 1000
Mbps
Cost Configured      : 1                                 Cost Calculated      : 4
State/Type           : DR                                Router Priority      : 1
DR                   : 10.10.10.2                        BDR                  :
10.10.10.1
Link LSAs            : 1                                 Checksum Sum         : 9168
AH Authentication    : SHA1, SPI 256 (Area)              Passive              : Yes

Codes: DR - Designated router  BDR - Backup Designated router

Interface vlan10 is Up, Line Protocol is Up, Instance ID 10
-------------------------------------------
VRF                  : default                           Process              : 1
IPv6 address         : fe80::7272:cfff:fe70:67ae         Area                 :
0.0.0.1
Status               : Up                                Network Type         :
Broadcast
Hello Interval       : 10    sec                         Dead Interval        : 40
 sec
Transit Delay        : 1     sec                         Retransmit Interval  : 5
 sec
BFD                  : Disabled                          Link Speed           : 1000
Mbps
Cost Configured      : 1                                 Cost Calculated      : 4
State/Type           : DR                                Router Priority      : 1
DR                   : 10.10.10.2                        BDR                  :
10.10.10.1
Link LSAs            : 1                                 Checksum Sum         : 9168
AH Authentication    : MD5, SPI 257                      Passive              : No

Codes: DR - Designated router  BDR - Backup Designated router

Interface vlan20 is Up, Line Protocol is Up, Instance ID 10
-------------------------------------------
VRF                  : default                           Process              : 1
IPv6 address         : fe80::7272:cfff:fe70:67ae         Area                 :
0.0.0.1
Status               : Up                                Network Type         :
Broadcast
Hello Interval       : 10    sec                         Dead Interval        : 40
 sec
Transit Delay        : 1     sec                         Retransmit Interval  : 5
 sec
BFD                  : Disabled                          Link Speed           : 1000
Mbps
Cost Configured      : 1                                 Cost Calculated      : 4
State/Type           : DR                                Router Priority      : 1
DR                   : 1.1.1.3                           BDR                  : NA
Link LSAs            : 1                                 Checksum Sum         : 9168
AH Authentication    : MD5, SPI 257                      Passive              : No
```

Showing overview information for OSPFv3 AF enabled interfaces for all VRFs in brief:

```
switch# show ospfv3-af ipv6 interface brief all-vrfs
VRF : default                          Process : 1
=================================================

Total Number of Interfaces: 2

Flags: P - Passive  A - Active
```

```
Interface     Area                Cost      State              Status      Flags
--------------------------------------------------------------------------------
1/1/1         0.0.0.0             1         BDR                 Up          P
vlan10        0.0.0.1             1         DR                  Up          A

VRF : red                                Process : 1
=================================================

Total Number of Interfaces: 1

Flags: P - Passive  A - Active


Interface     Area                Cost      State              Status      Flags
--------------------------------------------------------------------------------
1/1/2         0.0.0.1             1         BDR                 Up          P
tunnel1       0.0.0.0             100       DR                  Up          A
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ospfv3-af ipv6 lsdb

```
show ospfv3-af [<process-id>] ipv6 lsdb
   adv-router {<ROUTER-ID>|self}
   area <AREA-ID>
   lsid <link-state-id>
   all-vrfs|vrf <VRF-NAME>}
      as-external
      asbr-summary
      database-summary
      inter-area-prefix
      inter-area-router
      intra-area-prefix
      link
      network
      nssa-external
      router
```

```
summary
```

Description

Shows the OSPFv3 AF link state database summary for different OSPF LSAs (Link State Advertisement). Use the parameters to get information for a particular LSA.

Parameter	Description
<PROCESS-ID>	Enter an OSPFv3 AF process ID to display general OSPF information for a particular OSPF process. Range: 1 to 65535.
adv-router {*<ROUTER-ID>*|self}	Select to display link states for a particular advertising router. Specify either a Router ID of the advertising router or specify **self** to show self-originated link states.
area *<AREA-ID>*	Select to display information filtered for the specified area in one of the following formats. OSPFv3 AF area identifier in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. OSPFv3 AF area identifier in decimal format. Value: 0 to 4294967295.
lsid *<LINK-STATE-ID>*	Select to display information filtered by link state identifier specified in IPv4 address format (A.B.C.D).
all-vrfs|{vrf <vrf-name>	Select **all-vrfs** to display general OSPF information for all VRFs, or use the **vrf *<VRF-NAME>*** option to display information for a specific VRF. Optionally select one of the following parameters to filter the link state database information.
as-external	Show external link states (LSA type 5)
database-summary	Select to display the count of each type of LSA and each area in the database. **NOTE:** The **database-summary** parameter does not support the **area <area-id>**, **lsid <link-state-id>** or **adv-router {<router-id>|self}** parameters.
inter-area-prefix	Show inter-area prefix link states (LSA type 3)
inter-area-router	Show inter-area router link states (LSA type 4)
intra-area-prefix	Show intra-area prefix link states (LSA type 9
link	Show link states (LSA type 8)
network	Show network LSAs (LSA type 2).
nssa-external	Show NSSA external link states (LSA type 7).
router	Show router LSAs (LSA type 1).

Examples

Showing OSPFv3 AF link state database (LSDB) general information:

```
switch# show ospfv3-af ipv6 lsdb

Address-family : IPv6 Unicast
------------------------------

OSPF Router with ID (50.50.50.50) (Process ID 1 VRF default)
============================================================

Router Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID         Link Count  Bits
-------------------------------------------------------------------------------
40.40.40.40     930     0x80000004  0x2ea1      0            3           None
50.50.50.50     935     0x80000002  0x8b52      0            1           E
60.60.60.60     943     0x800003c5  0x9854      0            2           None

Network Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Router Count
-------------------------------------------------------------------------
60.60.60.60     944     0x80000001  0x7179      1360007168  2
50.50.50.50     935     0x80000001  0x516a      19          1

Inter Area Prefix Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
-----------------------------------------------------------------
40.40.40.40     929     0x80000001  0x2498      131072      FEC0:3344::/32
50.50.50.50     928     0x80000001  0x5b2f      65536       111::/64

Inter Area Router Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Destination Router ID
-------------------------------------------------------------------------------
40.40.40.40     929     0x80000001  0x2498      1           33.33.33.33

AS External Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
-----------------------------------------------------------------
40.40.40.40     264     0x80000001  0x24cc4     1           10::/64
40.40.40.40     675     0x80000001  0x5b00f     2           11::/64

NSSA External Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
-----------------------------------------------------------------
3.3.3.3         264     0x80000001  0x24ac2     1           200::/64

Link-local Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Interface
-------------------------------------------------------------------
50.50.50.50     264     0x80000001  0x653c4     19          1/1/1

Intra Area Prefix Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------------
-------------
ADV Router      Age     Seq#        Checksum    LSID        Referenced LS Type
Referenced LSID
-------------------------------------------------------------------------------
-------------
```

```
50.50.50.50     263     0x80000001  0x1da34     1           0x2001              0
50.50.50.50     264     0x80000001  0x2a45d     1           0x2002              19
```

Showing AS external link states:

```
switch# show ospfv3-af ipv6 lsdb as-external

Address-family : IPv6 Unicast
------------------------------

OSPF Router with ID (60.60.60.60) (Process ID 1 VRF default)
============================================================

AS External Link State Advertisements (Area 0.0.0.0)
--------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
--------------------------------------------------------------------
40.40.40.40     264     0x80000001  0x24cc4     1           10::/64
40.40.40.40     675     0x80000001  0x5b00f     2           11::/64
```

Showing the LSDB database summary:

```
switch# show ospfv3-af ipv6 lsdb database-summary

Address-family : IPv6 Unicast
------------------------------

OSPF Router with ID (10.1.1.1) (Process ID 1 VRF default)
=========================================================

Area 0.0.0.0 database summary
--------------------------
LSA Type             Count
--------------------------
Router               2
Network              1
Inter Area Prefix    1
Inter Area Router    0
NSSA External        0
Link                 3
Intra Area Prefix    3
--------------------------
Total                10

Process 1 database summary
--------------------------
LSA Type             Count
--------------------------
Router               2
Network              1
Inter Area Prefix    1
Inter Area Router    0
AS External          2
NSSA External        0
Link                 3
Intra Area Prefix    3
--------------------------
Total                12
```

Showing inter-area prefix LSAs:

```
switch# show ospfv3-af ipv6 lsdb inter-area-prefix

Address-family : IPv6 Unicast
------------------------------

OSPF Router with ID (6.6.6.6) (Process ID 1 VRF default)
=========================================================

Inter Area Prefix Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
-------------------------------------------------------------------
40.40.40.40     929     0x80000001  0x2498      131072      FEC0:3344::/32
50.50.50.50     928     0x80000001  0x5b2f      65536       111::/64
```

Showing network LSAs:

```
switch# show ospfv3-af ipv6 lsdb network

Address-family : IPv6 Unicast
------------------------------

OSPF Router with ID (50.50.50.50) (Process ID 1 VRF default)
=========================================================

Network Link State Advertisements (Area 0.0.0.0)
-----------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Router Count
-----------------------------------------------------------------------
60.60.60.60     944     0x80000001  0x7179      1360007168  2
50.50.50.50     935     0x80000001  0x516a      19          1
```

Showing NSSA external link states:

```
switch# show ospfv3-af ipv6 lsdb nssa-external

Address-family : IPv6 Unicast
------------------------------

OSPF Router with ID (2.2.2.1) (Process ID 1 VRF default)
=========================================================

NSSA External Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Prefix
-------------------------------------------------------------------
3.3.3.3         264     0x80000001  0x24ac2     1           200::/64
```

Showing router LSAs:

```
switch# show ospfv3-af ipv6 lsdb router

Address-family : IPv6 Unicast
------------------------------

OSPF Router with ID (50.50.50.50) (Process ID 1 VRF default)
=========================================================
```

```
Ruter Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID         Link Count  Bits
-------------------------------------------------------------------------------
40.40.40.40     930     0x80000004  0x2ea1      0            3           None
50.50.50.50     935     0x80000002  0x8b52      0            1           E
60.60.60.60     943     0x800003c5  0x9854      0            2           None
```

Showing inter-area router LSAs:

```
switch# show ospfv3-af ipv6 lsdb inter-area-router

Address-family : IPv6 Unicast
------------------------------

OSPF Router with ID (60.60.60.60) (Process ID 1 VRF default)
============================================================

Inter Area Router Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Destination Router ID
-------------------------------------------------------------------------------
40.40.40.40     929     0x80000001  0x2498      1           33.33.33.33
```

Showing link states:

```
switch# show ospfv3-af ipv6 lsdb link

Address-family : IPv6 Unicast
------------------------------

OSPF Router with ID (2.2.2.1) (Process ID 1 VRF default)
========================================================

Link-local Link State Advertisements (Area 0.0.0.0)
---------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Interface
---------------------------------------------------------------------
50.50.50.50     264     0x80000001  0x653c4     19          1/1/1

Link-local Link State Advertisements (Area 0.0.0.1)
---------------------------------------------------------------------
ADV Router      Age     Seq#        Checksum    LSID        Interface
---------------------------------------------------------------------
50.50.50.50     264     0x80000001  0xad324     20          1/1/2
```

Showing lintra-area prefix LSAs:

```
switch# show ospfv3-af ipv6 lsdb intra-area-prefix

Address-family : IPv6 Unicast
------------------------------

OSPF Router with ID (0.0.0.0) (Process ID 1 VRF default)
========================================================
```

```
Intra Area Prefix Link State Advertisements (Area 0.0.0.0)
-------------------------------------------------------------------------------
-------------
ADV Router      Age     Seq#        Checksum    LSID        Referenced LS Type
Referenced LSID
-------------------------------------------------------------------------------
-------------
50.50.50.50     263     0x80000001  0x1da34     1           0x2001              0
50.50.50.50     264     0x80000001  0x2a45d     1           0x2002              19
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ospfv3-af ipv6 neighbors

```
show ospfv3-af [<PROCESS-ID>] ipv6 neighbors [<NEIGHBOR-ID>]
   [interface <INTERFACE-NAME>] | vlan <VLAN ID> | lag <LAG ID> | loopback <LOOPBACK ID>
   | tunnel<tunnel id>][detail | summary]
   [all-vrfs | vrf <VRF-NAME>]
```

Description

Shows information about OSPFv3 AF neighbors.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 AF process ID to show OSPFv3 AF neighbor information for the particular OSPFv3 AF process. Range: 1 to 65535.
`neighbors` *<NEIGHBOR-ID>*	Shows information about a particular neighbor, specified in IPv4 format (A.B.C.D).
`interface` *<INTERFACE-NAME>*	Shows neighbor information only for the specified interface.
`detail`	Shows detailed information for the neighbors.

Parameter	Description
`summary`	Shows summary information for the neighbors.
`lag <LAG ID>`	Specifies lag interface name.
`loopback <LOOPBACK ID>`	Specifies the vlan interface name.
`all-vrfs`	Shows neighbor information for all VRFs.
`vrf <VRF-NAME>`	Specifies the name of a VRF. Default: default.

Examples

Showing OSPFv3 AF neighbors information:

```
switch# show ospfv3-af ipv6 neighbors
Address-family : IPv6 Unicast
------------------------------

OSPFv3 AF Process ID 1 VRF default
==================================

Total Number of Neighbors: 1

Neighbor ID      Priority  State             Interface
-------------------------------------------------------
3.3.3.3          1         FULL/DR           1/1/1
  Neighbor address fe80::7272:cfff:fe79:7510
```

Showing OSPFv3 AF neighbors information for neighbor 3.3.3.3:

```
switch# show ospfv3-af ipv6 neighbors 3.3.3.3
Address-family : IPv6 Unicast
------------------------------

VRF : default                          Process : 1
---------------------------------------------------------------
Router-Id         : 3.3.3.3            Area              : 0.0.0.0
Interface         : 1/1/1              Address           :
fe80::7272:cfff:fe79:7510
State             : FULL               Neighbor Priority : 1
Dead Timer Due    : 00:00:36           Options           : 0x13
Time since last state change : 00h:14m:45s
```

Showing detail OSPFv3 AF neighbors information for neighbor 3.3.3.3:

```
switch# show ospfv3-af ipv6 neighbors 3.3.3.3 detail
Address-family : IPv6 Unicast
------------------------------

VRF : default                          Process : 1
---------------------------------------------------------------
Router-Id         : 3.3.3.3            Area              : 0.0.0.0
Interface         : 1/1/1              Address           :
fe80::7272:cfff:fe79:7510
State             : FULL               Neighbor Priority : 1
```

```
DR                : 3.3.3.3           BDR               : 1.1.1.3
Dead Timer Due    : 00:00:36          Options           : 0x13
Retransmission Queue Length  : 0
Time Since Last State Change : 00h:14m:45s
```

Showing OSPFv3 AF neighbors information for interface 1/1/1:

```
switch# show ospfv3-af ipv6 neighbors 3.3.3.3 interface 1/1/1

Address-family : IPv6 Unicast
-----------------------------

VRF : default                         Process : 1
-------------------------------------------------------------
Router-Id         : 3.3.3.3           Area              : 0.0.0.0
Interface         : 1/1/1             Address           :
fe80::7272:cfff:fe79:7510
State             : FULL              Neighbor Priority : 1
Dead Timer Due    : 00:00:36          Options           : 0x13
Time Since Last State Change : 00h:14m:45s
```

Showing summary OSPFv3 AF neighbors information for a specific neighbor for all VRFs:

```
switch# show ospfv3-af ipv6 neighbors 3.3.3.3 summary all-vrfs

Address-family : IPv6 Unicast
-----------------------------

OSPFv3 AF Process ID 1 VRF default, Neighbor Summary
====================================================

Interface    Down Attempt Init TwoWay ExStart Exchange Loading Full Total
-------------------------------------------------------------------------
1/1/1        0    0       0    0      0       0        0       1    1
Total        0    0       0    0      0       0        0       1    1

OSPFv3 AF Process ID 1 VRF red, Neighbor Summary
================================================

Interface    Down Attempt Init TwoWay ExStart Exchange Loading Full Total
-------------------------------------------------------------------------
1/1/2        0    0       0    0      0       0        0       1    1
Total        0    0       0    0      0       0        0       1    1
```

Showing OSPFv3 AF neighbors information for VRF red:

```
switch# show ospfv3-af ipv6 neighbors vrf red
Address-family : IPv6 Unicast
-----------------------------

OSPFv3 AF Process ID 2 VRF red
==============================

Total Number of Neighbors: 1

Neighbor ID      Priority  State             Interface
-----------------------------------------------------
```

```
4.4.4.4          1          FULL/DR           1/1/2
  Neighbor address fe80::7272:cfff:fe79:7510
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ospfv3-af ipv6 routes

```
show ospfv3-af [<PROCESS-ID>] ipv6 routes [<PREFIX/LENGTH>]
   [all-vrfs | vrf <VRF-NAME>]
```

Description

Shows the OSPFv3 AF routing table information.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 AF process ID that shows information from the OSPFv3 AF routing table for the particular OSPFv3 AF process. Range: 1 to 65535.
<PREFIX/LENGTH>	Specifies the IPv6 destination prefix showing information about a particular destination prefix. For example, 2010:bd9::/32.
`all-vrfs`	Select to show OSPFv3 AF routing table information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Examples

Showing OSPFv3 AF routing table information:

```
switch# show ospfv3-af ipv6 routes
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2
```

```
Address-family : IPv6 Unicast
------------------------------

OSPFv3 AF Process ID 1 VRF default, Routing Table
-------------------------------------------------

Total Number of OSPFv3 AF Routes : 3

111::/64           (i) area:0.0.0.0
     directly attached to interface 1/1/1, cost 1 distance 110
2000::/64          (i) area:0.0.0.0
     directly attached to interface tunnel1, cost 100 distance 110
fd00::/64          (i) area:0.0.0.1
     directly attached to interface vlan10, cost 1 distance 110
```

Showing OSPFv3 AF routing table information for VRF red:

```
switch# show ospfv3-af ipv6 routes vrf red

Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

Address-family : IPv6 Unicast
------------------------------

OSPFv3 AF Process ID 2 VRF red, Routing Table
---------------------------------------------

Total Number of OSPFv3 AF Routes : 1

222::/64           (i) area:0.0.0.1
     directly attached to interface 1/1/2, cost 1 distance 110
```

Showing OSPFv3 AF routing table information for destination prefix fd00::/64:

```
switch# show ospfv3-af 1 ipv6 routes fd00::/64
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2
Address-family : IPv6 Unicast
------------------------------

OSPFv3 AF Process ID 1 VRF default, Routing Table for prefixes fd00::/64
------------------------------------------------------------------------

Total Number of OSPFv3 AF Routes : 1

fd00::/64          (i) area:0.0.0.1
     directly attached to interface vlan10, cost 1 distance 110
```

Showing OSPFv3 AF routing table information for destination prefix 2000::/64:

```
switch# show ospfv3-af 1 ipv6 routes 2000::/64
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2
Address-family : IPv6 Unicast
------------------------------
```

```
OSPFv3 AF Process ID 1 VRF default, Routing Table for prefixes 2000::/64
------------------------------------------------------------------------

Total Number of OSPFv3 AF Routes : 1

2000::/64           (i) area:0.0.0.0
     directly attached to interface tunnel1, cost 100 distance 110
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ospfv3-af ipv6 statistics

```
show ospfv3-af ipv6 [<PROCESS-ID>] statistics
     [all-vrfs | vrf <VRF-NAME>]
```

Description

Shows OSPFv3 AF statistics.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 AF process ID that shows information on the OSPFv3 AF SPF statistics for the particular OSPFv3 AF process. Range: 1 to 65535.
`all-vrfs`	Select to show OSPFv3 AF SPF statistics information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Examples

Showing OSPFv3 AF statistics information:

```
switch# show ospfv3-af ipv6 statistics

Address-family : IPv6 Unicast
------------------------------

OSPFv3 AF Process ID 1 VRF default, Statistics (cleared 3h 2m 21s ago)
----------------------------------------------------------------------

Unknown Interface Drops         : 0
Unknown Virtual Interface Drops : 0
Bad IPv6 Header Length Drops    : 0
Wrong OSPFv3 Version Drops      : 0
Bad Source IPv6 Drops           : 0
Resource Failure Drops          : 0
Bad Header Length Drops         : 0
Total Drops                     : 0
```

Showing OSPFv3 AF statistics information for VRF red:

```
switch# show ospfv3-af 2 ipv6 statistics vrf red

Address-family : IPv6 Unicast
------------------------------

OSPFv3 AF Process ID 2 VRF red, Statistics (cleared 3h 2m 30s ago)
------------------------------------------------------------------

Unknown Interface Drops         : 0
Unknown Virtual Interface Drops : 0
Bad IPv6 Header Length Drops    : 0
Wrong OSPFv3 Version Drops      : 0
Bad Source IPv6 Drops           : 0
Resource Failure Drops          : 0
Bad Header Length Drops         : 0
Total Drops                     : 0
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ospfv3-af ipv6 statistics interface

```
show ospfv3-af ipv6 [<PROCESS-ID>] statistics interface [<INTERFACE-NAME>| vlan <VLAN-ID>
| lag <LAG-ID> | loopback <LOOPBACK-ID> | tunnel <TUNNEL-ID>]
   [all-vrfs | vrf <VRF-NAME>]
```

Description

Shows the OSPFv3 AF statistics for the enabled OSPFv3 AF interfaces.

Parameter	Description
<PROCESS-ID>	Specifies an OSPFv3 AF process ID to show OSPF-enabled interface statistics information on the specified OSPFv3 AF process. Range: 1 to 65535.
<INTERFACE-NAME>	Selects to show information only for the specified interface.
`vlan` *<VLAN-ID>*	Specifies the vlan interface name.
`lag` *<LAG ID>*	Specifies lag interface name.
`loopback` *<LOOPBACK ID>*	Specifies the vlan interface name.
`tunnel` *<TUNNEL-ID>*	Specifies the tunnel interface name.
`all-vrfs`	Select to show OSPF-enabled interface statistics information for all VRFs.
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Example

Showing enabled OSPFv3 AF interfaces information for interface 1/1/1:

```
switch# show ospfv3-af ipv6 statistics interface 1/1/1

Address-family : IPv6 Unicast
------------------------------

OSPFv3 AF Process ID 1 VRF default, Interface 1/1/1 Statistics  (cleared 3h 22m
18s ago)

================================================================================
====

Tx Hello packets      : 1213             Rx Hello packets      : 1213
Tx Hello bytes        : 48516            Rx Hello bytes        : 48516
Tx DD packets         : 3                Rx DD packets         : 2
Tx DD bytes           : 164              Rx DD bytes           : 136
Tx LS request packets : 1                Rx LS request packets : 1
Tx LS request bytes   : 64               Rx LS request bytes   : 64
Tx LS update packets  : 22               Rx LS update packets  : 22
Tx LS update bytes    : 1520             Rx LS update bytes    : 1844
Tx LS ack packets     : 21               Rx LS ack packets     : 20
Tx LS ack bytes       : 976              Rx LS ack bytes       : 800

Total IPsec packets processed : 10
Total IPsec bytes processed   : 6400
Total Number of State Changes : 3
```

```
Number of LSAs                : 2
LSA Checksum Sum              : 21758

Total OSPFv3-AF Packets Discarded: 0
-----------------------------------

Reason                             Packets Dropped
-------------------------------------------------
Invalid Type                      0
Invalid Length                    0
Invalid Version                   0
Bad or Unknown Source             0
Area Mismatch                     0
Self-originated                   0
Duplicate Router ID               0
Interface Standby                 0
Total Hello Packets Dropped       0
  Hello Interval Mismatch         0
  Dead Interval Mismatch          0
  Options Mismatch                0
  MTU Mismatch                    0
  Neighbor Ignored                0
Resource Failures                 0
Bad LSA Length                    0
Others                            0
IPsec Authentication Errors       0

Total LSAs Ignored : 0
-----------------------
Bad Type           : 0
Bad Length         : 0
Invalid Data       : 0
Invalid Checksum   : 0
```

Showing enabled OSPFv3 AF interfaces information for tunnel interface:

```
switch# show ospfv3-af ipv6 statistics interface tunnel1

Address-family : IPv6 Unicast
-----------------------------

OSPFv3 AF Process ID 1 VRF default, Interface tunnel1 Statistics (cleared 0h16m32s
ago)

================================================================================
===

Tx Hello packets      : 100             Rx Hello packets      : 75
Tx Hello bytes        : 3900            Rx Hello bytes        : 3000
Tx DD packets         : 2               Rx DD packets         : 2
Tx DD bytes           : 136             Rx DD bytes           : 116
Tx LS request packets : 1               Rx LS request packets : 1
Tx LS request bytes   : 52              Rx LS request bytes   : 64
Tx LS update packets  : 3               Rx LS update packets  : 5
Tx LS update bytes    : 388             Rx LS update bytes    : 464
Tx LS ack packets     : 4               Rx LS ack packets     : 3
Tx LS ack bytes       : 244             Rx LS ack bytes       : 208

Total IPsec packets processed : 0
Total IPsec bytes processed   : 0
Total Number of State Changes : 2
```

```
Number of LSAs                : 2
LSA Checksum Sum              : 76646

Total OSPFv3-AF Packets Discarded: 0
-----------------------------------

Reason                            Packets Dropped
------------------------------------------------
Invalid Type                      0
Invalid Length                    0
Invalid Version                   0
Bad or Unknown Source             0
Area Mismatch                     0
Self-originated                   0
Duplicate Router ID               0
Interface Standby                 0
Total Hello Packets Dropped       0
  Hello Interval Mismatch         0
  Dead Interval Mismatch          0
  Options Mismatch                0
  MTU Mismatch                    0
  Neighbor Ignored                0
Resource Failures                 0
Bad LSA Length                    0
Others                            0
IPsec Authentication Errors       0
IPsec ESP Errors                  0

Total LSAs Ignored : 0
-----------------------
Bad Type           : 0
Bad Length         : 0
Invalid Data       : 0
Invalid Checksum   : 0
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ospfv3-af ipv6 virtual-links

`show ospfv3-af [<PROCESS-ID>] ipv6 virtual-links [brief]`

```
[all-vrfs | vrf <VRF-NAME>]
```

Description

Displays the current state and parameters of the OSPFv3 AF virtual links.

Parameter	Description
<PROCESS-ID>	Enter an OSPFv3 AF process ID to display information on the OSPFv3 AF virtual links for the particular OSPFv3 AF process. Range: 1 to 65535.
`brief`	Select to display brief overview information for the OSPFv3 AF virtual links.
`all-vrfs`	Select to display OSPFv3 AF virtual links information for all VRFs.
`vrf <VRF-NAME>`	Specifies the name of a VRF. Default: default.

Examples

Show OSPFv3 AF virtual links information:

```
switch# show ospfv3-af ipv6 virtual-links

Address-family : IPv6 Unicast
------------------------------

Virtual link to router 40.40.40.40 is up
----------------------------------------

VRF                 : default                       Process             : 21
Transit area        : 0.0.0.1                       Authentication      : No
Hello Interval      : 10                            Dead Interval       : 40
Transit Delay       : 1                             Retransmit Interval : 5
Number of Link LSAs : 0                             Checksum Sum        : 0
Authentication      : No
Number of State Changes : 4
```

Show OSPFv3 AF virtual links information with IPsec AH authentication:

```
switch# show ospfv3-af ipv6 virtual-links

Address-family : IPv6 Unicast
------------------------------

Virtual link to router 40.40.40.40 is up
----------------------------------------

VRF                 : default                       Process             : 21
Transit area        : 0.0.0.1                       Authentication      : No
Hello Interval      : 10                            Dead Interval       : 40
Transit Delay       : 1                             Retransmit Interval : 5
Number of Link LSAs : 0                             Checksum Sum        : 0
AH Authentication   : MD5  SPI: 1234
Number of State Changes : 4
```

Show OSPFv3 AF virtual links information with IPsec ESP authentication:

```
switch# show ospfv3-af ipv6 virtual-links

Address-family : IPv6 Unicast
------------------------------

Virtual link to router 40.40.40.40 is up
----------------------------------------

VRF                 : default                          Process             : 21
Transit area        : 0.0.0.1                          Authentication      : No
Hello Interval      : 10                               Dead Interval       : 40
Transit Delay       : 1                                Retransmit Interval : 5
Number of Link LSAs : 0                                Checksum Sum        : 0
ESP Authentication  : SHA1 Encryption: NULL  SPI: 1234
Number of State Changes : 4
```

Show brief overview information for OSPFv3 AF virtual links:

```
switch# show ospfv3-af ipv6 virtual-links brief

Address-family : IPv6 Unicast
------------------------------

OSPFv3 AF Process ID 1 VRF default
==================================

Total Number of Virtual Links: 1

Neighbor-ID     Transit-Area    Status
----------------------------------------
2.1.1.2         0.0.0.1         down
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show running-config

```
show running-config
```

Description

This command shows the configurations for all the protocols along with OSPFv3 AF.

Example

```
switch# show running-config
Current configuration:
!
!Version ArubaOS-CX generic x86-64.01.01.000X
!Schema version 0.1.8
vrf blue
vrf green
vrf red
led locator on
!
!
!
!
!
!
!
vlan 1,10
interface 1/1/1
   no shutdown
   ip address 1.1.1.1/24
   ipv6 address 111::1/64
   ip ospf 1 area 0.0.0.0
   ospfv3-af 1 ipv4 area 0.0.0.0
   ospfv3-af ipv4 authentication ipsec spi 300 sha1 ciphertext
AQBapX1/Ow3jK2zC6WCOv3mGwmvDQ02L9WKRTxVPtSyXpUNqBQAAACoBmj2c
interface 1/1/2
   no shutdown
   vrf attach red
   ip address 2.2.2.2/24
   ipv6 address 222::2/64
   ip ospf 2 area 0.0.0.1
   ospfv3-af 1 ipv4 area 0.0.0.1
   ospfv3-af ipv4 authentication ipsec spi 300 md5 ciphertext
AQBapYH+i/6HV2w+yyJCb45XhHPFfLNIt876Mh9EMRyf0XgwBAAAAGE/tKs=
interface 1/1/3
   no shutdown
   ip address 3.3.3.1/24
   ipv6 address 333::3/64
   ospfv3-af 1 ipv6 area 0.0.0.0 instance 1
   no ospfv3-af ipv6 passive
   ospfv3-af 1 ipv4 area 0.0.0.0 instance 65
interface 1/1/4
   no shutdown
   ipv6 address 1000::1/64
interface tunnel 1 mode ip 6in6
   no shutdown
   source ipv6 1000::1
   destination ipv6 1000::2
   ipv6 address 2000::1/64
   ospfv3-af 1 ipv4 area 0.0.0.0
   ospfv3-af ipv4 passive
   ospfv3-af ipv4 hello-interval 22
   ospfv3-af ipv4 dead-interval 88
   ospfv3-af ipv4 transit-delay 44
   ospfv3-af ipv4 priority 88
   ospfv3-af ipv4 cost 8
   ospfv3-af ipv4 network point-to-point
```

```
   ospfv3-af ipv4 shutdown
interface vlan10
   ip address 10.1.1.1/24
   ipv6 address fd00::1/64
   ip ospf 1 area 0.0.0.0
   ospfv3-af 1 ipv4 area 0.0.0.1
   ospfv3-af ipv4 encryption ipsec spi 401 sha1 ciphertext
AQBapd+HuD8Y4kMJkh5dFzG6OTBvW4iY+ueK4xpETxZPsZ1FDgAAAE3pkLmvMCtdFk2zz/LS
aes ciphertext
AQBapbPnk0Kl9giahf2dWGRHcU+6mTFiWzfxeWUXuWGYQ4teIgAAAOG7kcVrkrKV754nX+Li1ZDYPHG5xs
apOnyZJMv7cfHEctA=
   ospfv3-af 1 ipv6 area 0.0.0.1
   ospfv3-af ipv4 encryption ipsec spi 401 sha1 ciphertext
AQBapd+HuD8Y4kMJkh5dFzG6OTBvW4iY+ueK4xpETxZPsZ1FDgAAAE3pkLmvMCtdFk2zz/LS
   aes ciphertext
AQBapbPnk0Kl9giahf2dWGRHcU+6mTFiWzfxeWUXuWGYQ4teIgAAAOG7kcVrkrKV754nX+Li1ZDYPHG5xs
apOnyZJMv7cfHEctA=
interface loopback 2
   ipv6 address 99::1/64
   ospfv3-af 1 ipv4 area 0.0.0.0
router ospfv3-af 1
   router-id 1.1.1.1
   reference-bandwidth 4000
   address-family ipv4 unicast
        router-id 2.2.2.1
        area 0.0.0.0
        area 0.0.0.1
        area 0.0.0.2
        area 0.0.0.3
        distribute-list prefix list2 in
        redistribute static
        reference-bandwidth 4000
        active-backbone stub-default-route
    exit-address-family
    address-family ipv6 unicast
        router-id 2.2.2.2
        area 0.0.0.0
        area 0.0.0.1
        area 0.0.0.2
        area 0.0.0.3
        distribute-list prefix list2 in
        passive-interface default
        redistribute static
        reference-bandwidth 4000
        active-backbone stub-default-route
        area 0.0.0.1 virtual-link 1.1.1.1
                hello-interval 15
                dead-interval 60
                authentication ipsec spi 256 md5 ciphertext
AQBapYH+i/6HV2w+yyJCb45XhHPFfLNIt876Mh9EMRyf0XgwBAAAAGE/tKs=
                exit
        area 0.0.0.1 authentication ipsec spi 256 md5 ciphertext
AQBapYH+i/6HV2w+yyJCb45XhHPFfLNIt876Mh9EMRyf0XgwBAAAAGE/tKs=
        area 0.0.0.2 virtual-link 2.2.2.2
                encryption ipsec spi 3456 md5 ciphertext
AQBapVuqYeTJuT8/BoIlC2z4MxuT3VVmiduKtDwIP1jbCCxZCAAAAEe+Xod/lZX8 null
                exit
        area 0.0.0.3 virtual-link 3.3.3.3
                encryption ipsec spi 5678 sha1 ciphertext
AQBapVuqYeTJuT8/BoIlC2z4MxuT3VVmiduKtDwIP1jbCCxZCAAAAEe+Xod/lZX8 des
ciphertext AQBapVuqYeTJuT8/BoIlC2z4MxuT3VVmiduKtDwIP1jbCCxZCAAAAEe+Xod/lZX8
                exit
```

```
        exit-address-family
router ospfv3-af 2
    router-id 1.1.1.3
    distance 88
    max-metric router-lsa
    throttle spf start-time 100 hold-time 1000 max-wait-time 10000
    address-family ipv4 unicast
        router-id 1.1.1.3
        distance 88
        area 0.0.0.0
        area 0.0.0.1
        passive-interface default
        redistribute ospfv3-af 1 route-map ospfv3_af4_routes
        max-metric router-lsa
        throttle spf start-time 100 hold-time 1000 max-wait-time 10000
        area 0.0.0.1 encryption ipsec spi 400 md5 ciphertext
AQBapfwg9p/z+rRS1mhk9jXzbOFBd82j2OSvEbAbI2f4KL+ADAAAAL3hhFDJlUbBwu1fvA==
des ciphertext
AQBapfTQ8jUL+zbNyX9ZKOF8Lc5iR6dF543pwK2egyCWSWkkEgAAAE0etJaie2d/i2XF8pKkRpQvpw==
        exit-address-family
    address-family ipv6 unicast
        router-id 1.1.1.4
        distance 88
        area 0.0.0.0
        area 0.0.0.1
        passive-interface default
        redistribute ospfv3-af 1 route-map ospfv3_af6_routes
        max-metric router-lsa
        throttle spf start-time 100 hold-time 1000 max-wait-time 10000
        area 0.0.0.1 encryption ipsec spi 400 md5 ciphertext
AQBapfwg9p/z+rRS1mhk9jXzbOFBd82j2OSvEbAbI2f4KL+ADAAAAL3hhFDJlUbBwu1fvA==
des ciphertext
AQBapfTQ8jUL+zbNyX9ZKOF8Lc5iR6dF543pwK2egyCWSWkkEgAAAE0etJaie2d/i2XF8pKkRpQvpw==
        exit-address-family
router ospfv3-af 1 vrf red
    reference-bandwidth 40000
    graceful-restart restart-interval 88
    no graceful-restart helper
    distance intra-area 34 inter-area 66 external 3
    address-family ipv4 unicast
        area 0.0.0.1
        redistribute ospf 2
        redistribute ospfv3-af 2
        reference-bandwidth 40000
        graceful-restart restart-interval 88
        no graceful-restart helper
        distance intra-area 34 inter-area 66 external 3
    exit-address-family
    address-family ipv6 unicast
        area 0.0.0.1
        redistribute ospfv3 2
        redistribute ospfv3-af 2
        reference-bandwidth 40000
        graceful-restart restart-interval 88
    max-metric router-lsa include-stub
    graceful-restart restart-interval 88
    no graceful-restart helper
    graceful-restart ignore-lost-interface
    distance intra-area 34 inter-area 66
    address-family ipv4 unicast
        router-id 1.1.1.8
        area 0.0.0.1
```

```
        max-metric router-lsa include-stub
        graceful-restart restart-interval 88
        no graceful-restart helper
        graceful-restart ignore-lost-interface
        distance intra-area 34 inter-area 66
    exit-address-family
    address-family ipv6 unicast
        router-id 1.1.1.9
        area 0.0.0.1
        graceful-restart restart-interval 88
        no graceful-restart helper strict-lsa-check
        max-metric router-lsa include-stub
        distance intra-area 34 inter-area 66
    exit-address-family
router ospfv3-af 1 vrf blue
    graceful-restart restart-interval 99
    no graceful-restart helper
    max-metric router-lsa on-startup 100
    address-family ipv4 unicast
        graceful-restart restart-interval 99
        no graceful-restart helper
        area 0.0.0.0
        max-metric router-lsa on-startup 100
    exit-address-family
    address-family ipv6 unicast
        graceful-restart restart-interval 99
        no graceful-restart helper
        area 0.0.0.0
        max-metric router-lsa on-startup 100
    exit-address-family
router ospfv3-af 1 vrf green
    max-metric router-lsa include-stub on-startup 100
    address-family ipv4 unicast
        area 0.0.0.0
        max-metric router-lsa include-stub on-startup 100
    exit-address-family
    address-family ipv6 unicast
        area 0.0.0.0
        max-metric router-lsa include-stub on-startup 100
    exit-address-family
switch#
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P	Operator (>) or Manager (#)	

Platforms	Command context	Authority
8360 9300 9300S 10000		

show running-config all

```
show running-config all
```

Description

Displays the current running configuration, including default values for all protocols such as OSPFv3 AF.

Examples

Showing running-config without default values:

```
switch# show running-config
Current configuration:
!
!Version ArubaOS-CX Virtual.10.05.0001BC-26-g94067ba
!export-password: default
led locator on
vrf blue
vrf green
vrf red
!
!
!
!
ssh server vrf mgmt
vlan 1
interface mgmt
    no shutdown
    ip dhcp
interface 1/1/1
    no shutdown
    ip address 1.1.1.1/24
    ipv6 address 111::1/64
    ipv6 address 99::1/64
    ospfv3-af 1 ipv4 area 0.0.0.0
interface 1/1/2
    no shutdown
interface loopback 2
    ipv6 address 929::1/64
    ospfv3-af 1 ipv4 area 0.0.0.0
interface loopback 3
    ipv6 address 123::1/64
    ospfv3-af 1 ipv4 area 0.0.0.0
    ospfv3-af 1 ipv6 area 0.0.0.0
!
!
!
!
!
router ospfv3-af 1 vrf blue
    max-metric router-lsa on-startup 100
    address-family ipv4 unicast
        distribute-list prefix list2 in
```

```
        max-metric router-lsa on-startup 100
        active-backbone stub-default-route
        area 0.0.0.0
    exit-address-family
    address-family ipv6 unicast
        distribute-list prefix list2 in
        max-metric router-lsa on-startup 100
        active-backbone stub-default-route
        area 0.0.0.0
    exit-address-family
router ospfv3-af 1
    router-id 1.1.1.1
    reference-bandwidth 4000
    address-family ipv4 unicast
        router-id 1.1.1.2
        reference-bandwidth 4000
        redistribute host-routes
        redistribute static
        area 0.0.0.0
        area 0.0.0.1
    exit-address-family
    address-family ipv6 unicast
        router-id 1.1.1.3
        reference-bandwidth 4000
        redistribute host-routes
        redistribute static
        area 0.0.0.0
        area 0.0.0.1
        area 0.0.0.1 virtual-link 1.1.1.2
            dead-interval 60
            hello-interval 15
            authentication ipsec spi 256 md5 ciphertext
            AQBapYH+i/6HV2w+yyJCb45XhHPFfLNIt876Mh9EMRyf0XgwBAAAAGE/tKs=
            exit
        area 0.0.0.2
        area 0.0.0.2 virtual-link 2.2.2.2
            encryption ipsec spi 3456 md5 ciphertext
            AQBapVuqYeTJuT8/BoIlC2z4MxuT3VVmiduKtDwIP1jbCCxZCAAAAEe+Xod/lZX8 null
            exit
        area 0.0.0.3
        area 0.0.0.3 virtual-link 3.3.3.3
            encryption ipsec spi 5678 sha1 ciphertext
            AQBapVuqYeTJuT8/BoIlC2z4MxuT3VVmiduKtDwIP1jbCCxZCAAAAEe+Xod/lZX8 des
            ciphertext AQBapVuqYeTJuT8/BoIlC2z4MxuT3VVmiduKtDwIP1jbCCxZCAAAAEe+
            Xod/lZX8
            exit
    exit-address-family
router ospfv3-af 2
    router-id 1.1.1.3
    max-metric router-lsa
    address-family ipv4 unicast
        router-id 1.1.1.4
        max-metric router-lsa
        area 0.0.0.0
        area 0.0.0.1
        area 0.0.0.1 encryption ipsec spi 400 md5 ciphertext

AQBapfwg9p/z+rRS1mhk9jXzbOFBd82j2OSvEbAbI2f4KL+ADAAAAL3hhFDJlUbBwu1fvA==des
        ciphertext
AQBapfTQ8jUL+zbNyX9ZKOF8Lc5iR6dF543pwK2egyCWSWkkEgAAAE0etJaie2d/
        i2XF8pKkRpQvpw==
    exit-address-family
```

```
    address-family ipv6 unicast
        router-id 1.1.1.5
        max-metric router-lsa
        area 0.0.0.0
        area 0.0.0.1
        area 0.0.0.1 encryption ipsec spi 400 md5 ciphertext

AQBapfwg9p/z+rRS1mhk9jXzbOFBd82j2OSvEbAbI2f4KL+ADAAAAL3hhFDJlUbBwu1fvA==des
        ciphertext
AQBapfTQ8jUL+zbNyX9ZKOF8Lc5iR6dF543pwK2egyCWSWkkEgAAAE0etJaie2d/
        i2XF8pKkRpQvpw==
    exit-address-family
router ospfv3-af 1 vrf green
    max-metric router-lsa include-stub on-startup 100
    address-family ipv4 unicast
        max-metric router-lsa include-stub on-startup 100
        area 0.0.0.0
    exit-address-family
    address-family ipv6 unicast
        max-metric router-lsa include-stub on-startup 100
        area 0.0.0.0
    exit-address-family
router ospfv3-af 1 vrf red
    reference-bandwidth 40000
    address-family ipv4 unicast
        reference-bandwidth 40000
        redistribute ospfv3 2
        area 0.0.0.1
    exit-address-family
    address-family ipv6 unicast
        reference-bandwidth 40000
        redistribute ospfv3 2
        area 0.0.0.1
    exit-address-family
router ospfv3-af 2 vrf red
    router-id 1.1.1.7
    max-metric router-lsa include-stub
    address-family ipv4 unicast
        router-id 1.1.1.7
        max-metric router-lsa include-stub
        area 0.0.0.1
    exit-address-family
    address-family ipv6 unicast
        router-id 1.1.1.7
        max-metric router-lsa include-stub
        area 0.0.0.1
    exit-address-family
https-server vrf mgmt
```

Showing running-config with default values:

```
switch# show running-config all
Current configuration:
!
!
!
interface 1/1/1
    no shutdown
    mtu 1500
    qos trust none
    apply qos schedule-profile factory-default
```

```
    lldp transmit
    lldp receive
    lldp med poe
    no lldp med poe priority-override
    lldp dot3 poe
    lldp dot3 macphy
    lldp med network-policy
    lldp med capability
    lldp med topology-trap
    cdp
    sflow
    routing
    vrf attach default
    ip mtu 1500
    ip address 1.1.1.1/24
    ipv6 address 111::1/64 tag 0
    ipv6 address 99::1/64 tag 0
    no ip urpf-check
    ipv6 nd mtu 1500
    no ip proxy-arp
    ipv6 nd dad attempts 1
    ipv6 nd hop-limit 64
    ipv6 nd ns-interval 1000
    ipv6 nd ra min-interval 200
    ipv6 nd ra max-interval 600
    ipv6 nd ra lifetime 1800
    ipv6 nd ra reachable-time 0
    no ipv6 address autoconfig
    ipv6 nd ra retrans-timer 0
    no ipv6 nd ra managed-config-flag
    no ipv6 nd ra other-config-flag
    ipv6 nd suppress-ra
    ipv6 nd router-preference medium
    no ip irdp
    ip irdp multicast
    ip irdp minadvertinterval 450
    ip irdp maxadvertinterval 600
    ip irdp holdtime 1800
    ip irdp preference 0
    arp timeout 1800
    ipv6 neighbor timeout 1800
    power-over-ethernet
    no power-over-ethernet pre-std-detect
    power-over-ethernet priority low
    power-over-ethernet allocate-by usage
    power-over-ethernet assigned-class 4
    ospfv3-af 1 ipv4 area 0.0.0.0
    ospfv3-af ipv4 dead-interval 40
    ospfv3-af ipv4 hello-interval 10
    ospfv3-af ipv4 retransmit-interval 5
    ospfv3-af ipv4 transit-delay 1
    ospfv3-af ipv4 priority 1
    ospfv3-af ipv4 cost 4
    ospfv3-af ipv4 network broadcast
!
!
!
interface loopback 2
    vrf attach default
    ipv6 address 929::1/64 tag 0
    no ip urpf-check
    arp timeout 1800
```

```
    ipv6 neighbor timeout 1800
    ospfv3-af 1 ipv4 area 0.0.0.0
    ospfv3-af ipv4 dead-interval 40
    ospfv3-af ipv4 hello-interval 10
    ospfv3-af ipv4 retransmit-interval 5
    ospfv3-af ipv4 transit-delay 1
    ospfv3-af ipv4 priority 1
    ospfv3-af ipv4 network broadcast
!
!
!
interface loopback 3
    vrf attach default
    ipv6 address 123::1/64 tag 0
    no ip urpf-check
    arp timeout 1800
    ipv6 neighbor timeout 1800
    ospfv3-af 1 ipv4 area 0.0.0.0
    ospfv3-af ipv4 dead-interval 40
    ospfv3-af ipv4 hello-interval 10
    ospfv3-af ipv4 retransmit-interval 5
    ospfv3-af ipv4 transit-delay 1
    ospfv3-af ipv4 priority 1
    ospfv3-af ipv4 network broadcast
    ospfv3-af 1 ipv6 area 0.0.0.0
    ospfv3-af ipv6 dead-interval 40
    ospfv3-af ipv6 hello-interval 10
    ospfv3-af ipv6 retransmit-interval 5
    ospfv3-af ipv6 transit-delay 1
    ospfv3-af ipv6 priority 1
    ospfv3-af ipv6 network broadcast
!
!
!
router ospfv3-af 1 vrf blue
    enable
    reference-bandwidth 100000
    max-metric router-lsa on-startup 100
    timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
    timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
    timers lsa-arrival 1000
    graceful-restart restart-interval 120
    distance 110
    address-family ipv4 unicast
        enable
        reference-bandwidth 100000
        max-metric router-lsa on-startup 100
        maximum-paths 4
        timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
        timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
        timers lsa-arrival 1000
        graceful-restart restart-interval 120
        default-metric 25
        distance 110
        active-backbone stub-default-route
        area 0.0.0.0
        area 0.0.0.0 default-metric 1
    exit-address-family
    address-family ipv6 unicast
        enable
        reference-bandwidth 100000
        max-metric router-lsa on-startup 100
```

```
        maximum-paths 4
        timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
        timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
        timers lsa-arrival 1000
        graceful-restart restart-interval 120
        default-metric 25
        distance 110
        active-backbone stub-default-route
        area 0.0.0.0
        area 0.0.0.0 default-metric 1
    exit-address-family
router ospfv3-af 1 vrf default
    enable
    router-id 1.1.1.1
    reference-bandwidth 4000
    timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
    timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
    timers lsa-arrival 1000
    graceful-restart restart-interval 120
    distance 110
    address-family ipv4 unicast
        enable
        router-id 1.1.1.1
        reference-bandwidth 4000
        maximum-paths 4
        timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
        timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
        timers lsa-arrival 1000
        graceful-restart restart-interval 120
        default-metric 25
        distance 110
        redistribute static
        no active-backbone stub-default-route
        area 0.0.0.0
        area 0.0.0.0 default-metric 1
        area 0.0.0.1
        area 0.0.0.1 default-metric 1
        area 0.0.0.2
        area 0.0.0.2 default-metric 1
        area 0.0.0.3
        area 0.0.0.3 default-metric 1
    exit-address-family
    address-family ipv6 unicast
        enable
        router-id 1.1.1.1
        reference-bandwidth 4000
        maximum-paths 4
        timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
        timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
        timers lsa-arrival 1000
        graceful-restart restart-interval 120
        default-metric 25
        distance 110
        redistribute static
        no active-backbone stub-default-route
        area 0.0.0.0
        area 0.0.0.0 default-metric 1
        area 0.0.0.1
        area 0.0.0.1 default-metric 1
        area 0.0.0.1 virtual-link 1.1.1.2
            dead-interval 60
            hello-interval 15
```

```
            retransmit-interval 5
            transit-delay 1
            authentication ipsec spi 256 md5 ciphertext
            AQBapYH+i/6HV2w+yyJCb45XhHPFfLNIt876Mh9EMRyf0XgwBAAAAGE/tKs=
            exit
        area 0.0.0.2
        area 0.0.0.2 default-metric 1
        area 0.0.0.2 virtual-link 2.2.2.2
            dead-interval 40
            hello-interval 10
            retransmit-interval 5
            transit-delay 1
            encryption ipsec spi 3456 md5 ciphertext
            AQBapVuqYeTJuT8/BoIlC2z4MxuT3VVmiduKtDwIP1jbCCxZCAAAAEe+Xod/lZX8 null
            exit
        area 0.0.0.3
        area 0.0.0.3 default-metric 1
        area 0.0.0.3 virtual-link 3.3.3.3
        dead-interval 40
        hello-interval 10
        retransmit-interval 5
        transit-delay 1
        encryption ipsec spi 5678 sha1 ciphertext
        AQBapVuqYeTJuT8/BoIlC2z4MxuT3VVmiduKtDwIP1jbCCxZCAAAAEe+Xod/lZX8 des
        ciphertext
AQBapVuqYeTJuT8/BoIlC2z4MxuT3VVmiduKtDwIP1jbCCxZCAAAAEe+Xod/lZX8
        exit
    exit-address-family
router ospfv3-af 2 vrf default
    enable
    router-id 1.1.1.3
    reference-bandwidth 100000
    max-metric router-lsa
    timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
    timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
    timers lsa-arrival 1000
    graceful-restart restart-interval 120
    distance 110
    address-family ipv4 unicast
        enable
        router-id 1.1.1.3
        reference-bandwidth 100000
        max-metric router-lsa
        maximum-paths 4
        timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
        timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
        timers lsa-arrival 1000
        graceful-restart restart-interval 120
        default-metric 25
        distance 110
        no active-backbone stub-default-route
        area 0.0.0.0
        area 0.0.0.0 default-metric 1
        area 0.0.0.1
        area 0.0.0.1 encryption ipsec spi 400 md5 ciphertext

AQBapfwg9p/z+rRS1mhk9jXzbOFBd82j2OSvEbAbI2f4KL+ADAAAAL3hhFDJlUbBwu1fvA==des
        ciphertext
AQBapfTQ8jUL+zbNyX9ZKOF8Lc5iR6dF543pwK2egyCWSWkkEgAAAE0etJaie2d/
        i2XF8pKkRpQvpw==
        area 0.0.0.1 default-metric 1
    exit-address-family
```

```
    address-family ipv6 unicast
        enable
        router-id 1.1.1.3
        reference-bandwidth 100000
        max-metric router-lsa
        maximum-paths 4
        timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
        timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
        timers lsa-arrival 1000
        graceful-restart restart-interval 120
        default-metric 25
        distance 110
        no active-backbone stub-default-route
        area 0.0.0.0
        area 0.0.0.0 default-metric 1
        area 0.0.0.1
        area 0.0.0.1 encryption ipsec spi 400 md5 ciphertext

AQBapfwg9p/z+rRS1mhk9jXzbOFBd82j2OSvEbAbI2f4KL+ADAAAAL3hhFDJlUbBwu1fvA==des
        ciphertext
AQBapfTQ8jUL+zbNyX9ZKOF8Lc5iR6dF543pwK2egyCWSWkkEgAAAE0etJaie2d/
        i2XF8pKkRpQvpw==
        area 0.0.0.1 default-metric 1
    exit-address-family
router ospfv3-af 1 vrf green
    enable
    reference-bandwidth 100000
    max-metric router-lsa include-stub on-startup 100
    timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
    timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
    timers lsa-arrival 1000
    graceful-restart restart-interval 120
    distance 110
    address-family ipv4 unicast
        enable
        reference-bandwidth 100000
        max-metric router-lsa include-stub on-startup 100
        maximum-paths 4
        timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
        timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
        timers lsa-arrival 1000
        graceful-restart restart-interval 120
        default-metric 25
        distance 110
        no active-backbone stub-default-route
        area 0.0.0.0
        area 0.0.0.0 default-metric 1
    exit-address-family
    address-family ipv6 unicast
        enable
        reference-bandwidth 100000
        max-metric router-lsa include-stub on-startup 100
        maximum-paths 4
        timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
        timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
        timers lsa-arrival 1000
        graceful-restart restart-interval 120
        default-metric 25
        distance 110
        no active-backbone stub-default-route
        area 0.0.0.0
        area 0.0.0.0 default-metric 1
```

```
    exit-address-family
router ospfv3-af 1 vrf red
    enable
    reference-bandwidth 40000
    timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
    timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
    timers lsa-arrival 1000
    graceful-restart restart-interval 120
    distance 110
    address-family ipv4 unicast
        enable
        reference-bandwidth 40000
        maximum-paths 4
        timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
        timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
        timers lsa-arrival 1000
        graceful-restart restart-interval 120
        default-metric 25
        distance 110
        redistribute ospfv3 2
        no active-backbone stub-default-route
        area 0.0.0.1
        area 0.0.0.1 default-metric 1
    exit-address-family
    address-family ipv6 unicast
        enable
        reference-bandwidth 40000
        maximum-paths 4
        timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
        timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
        timers lsa-arrival 1000
        graceful-restart restart-interval 120
        default-metric 25
        distance 110
        redistribute ospfv3 2
        no active-backbone stub-default-route
        area 0.0.0.1
        area 0.0.0.1 default-metric 1
    exit-address-family
router ospfv3-af 2 vrf red
        enable
        router-id 1.1.1.7
        reference-bandwidth 100000
        max-metric router-lsa include-stub
        timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
        timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
        timers lsa-arrival 1000
        graceful-restart restart-interval 120
    distance 110
    address-family ipv4 unicast
        enable
        router-id 1.1.1.7
        reference-bandwidth 100000
        max-metric router-lsa include-stub
        maximum-paths 4
        timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
        timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
        timers lsa-arrival 1000
        graceful-restart restart-interval 120
        default-metric 25
        distance 110
        no active-backbone stub-default-route
```

```
        area 0.0.0.1
        area 0.0.0.1 default-metric 1
    exit-address-family
    address-family ipv6 unicast
        enable
        router-id 1.1.1.7
        reference-bandwidth 100000
        max-metric router-lsa include-stub
        maximum-paths 4
        timers throttle spf start-time 200 hold-time 1000 max-wait-time 5000
        timers throttle lsa start-time 5000 hold-time 0 max-wait-time 0
        timers lsa-arrival 1000
        graceful-restart restart-interval 120
        default-metric 25
        distance 110
        no active-backbone stub-default-route
        area 0.0.0.1
        area 0.0.0.1 default-metric 1
    exit-address-family
!
!
!
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show tech ospfv3-af

```
show tech ospfv3-af [local-file]
```

Description

Shows the output of the following OSPFv3 AF commands:

- **show ospfv3-af all vrfs**
- **show ospfv3-af interface all-vrfs**
- **show ospfv3-af neighbors detail all-vrfs**
- **show ospfv3-af neighbors summary all-vrfs**

- **show ospfv3-af route all-vrfs**
- **show ospfv3-af statistics all-vrfs**
- **show ospfv3-af statistics interface all-vrfs**
- **show ospfv3-af border-routers all-vrfs**
- **show ospfv3-af lsdb all-vrfs database-summary**
- **show ospfv3-af virtual-links all-vrfs**

Parameter	Description
`local-file`	Captures the command output into a local file. Optional.

Examples

```
switch(config-if)# show tech ospfv3-af
==================================================
Show Tech executed on Tue Sep 12 09:42:42 2017
==================================================
==================================================
[Begin] Feature ospfv3-af
==================================================


*********************************
Command : show ospfv3-af ipv6 neighbor summary all-vrfs
*********************************
Address-family : IPv6 Unicast
-----------------------------

OSPFv3 AF Process ID 1 VRF default, Neighbor Summary
================================================

Interface    Down Attempt Init TwoWay ExStart Exchange Loading Full Total
-------------------------------------------------------------------------
1/1/1        0    0       0    0      0       0        0       1    1
Total        0    0       0    0      0       0        0       1    1

OSPFv3 AF Process ID 1 VRF red, Neighbor Summary
============================================

Interface    Down Attempt Init TwoWay ExStart Exchange Loading Full Total
-------------------------------------------------------------------------
1/1/2        0    0       0    0      0       0        0       1    1
Total        0    0       0    0      0       0        0       1    1

*********************************
Command : show ospfv3-af ipv6 statistics all-vrfs
*********************************
Address-family : IPv6 Unicast
-----------------------------

OSPFv3 AF Process ID 1 VRF default, Statistics (cleared 0h 4m 46s ago)
----------------------------------------------------------------------

Unknown Interface Drops         : 0
Unknown Virtual Interface Drops : 0
Bad IPv6 Header Length Drops    : 0
Wrong OSPFv3 Version Drops      : 0
Bad Source IPv6 Drops           : 0
```

```
Resource Failure Drops          : 0
Bad Header Length Drops         : 0
Total Drops                     : 0

OSPFv3 AF Process ID 2 VRF red, Statistics (cleared 0h 4m 45s ago)
------------------------------------------------------------------

Unknown Interface Drops         : 0
Unknown Virtual Interface Drops : 0
Bad IPv6 Header Length Drops    : 0
Wrong OSPFv3 Version Drops      : 0
Bad Source IPv6 Drops           : 0
Resource Failure Drops          : 0
Bad Header Length Drops         : 0
Total Drops                     : 0

********************************
Command : show ospfv3-af ipv6 neighbor detail all-vrfs
********************************
Address-family : IPv6 Unicast
------------------------------

Neighbor 3.3.3.3, Interface address fe80::7272:cfff:fec6:6e6c
-------------------------------------------------------------

Process ID 1 VRF default, in Area 0.0.0.0 via Interface 1/1/1
    State is FULL, Neighbor priority is 1
    DR is 3.3.3.3, BDR is 1.1.1.3
    Options is 0x13
    Dead timer due in 00:00:34
    Retransmission queue length 0

Neighbor 4.4.4.4, Interface address fe80::7272:cfff:fec6:6e6c
-------------------------------------------------------------

Process ID 2 VRF red, in Area 0.0.0.1 via Interface 1/1/2
    State is FULL, Neighbor priority is 1
    DR is 4.4.4.4, BDR is 1.1.1.7
    Options is 0x13
    Dead timer due in 00:00:35
    Retransmission queue length 0


********************************
Command : show ospfv3-af ipv6 route all-vrfs
********************************
Address-family : IPv6 Unicast
------------------------------

Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPFv3 AF Process ID 1 VRF default, Routing Table
-------------------------------------------------

Total Number of OSPFv3 Routes : 2

111::/64           (i) area:0.0.0.0
     directly attached to interface 1/1/1, cost 1 distance 110
fd00::/64          (i) area:0.0.0.1
     directly attached to interface vlan10, cost 1 distance 110
```

```
OSPFv3 AF Process ID 2 VRF red, Routing Table
---------------------------------------------

Total Number of OSPFv3 Routes : 1

222::/64           (i) area:0.0.0.1
     directly attached to interface 1/1/2, cost 1 distance 110


*********************************
Command : show ospfv3-af ipv6 all-vrfs
*********************************
Address-family : IPv6 Unicast
------------------------------

Routing Process 1 with ID: 1.1.1.3 VRF default
----------------------------------------------

Graceful-restart is configured
Restart Interval: 120, State: inactive
Last Graceful Restart Exit Status: none
Maximum Paths to Destination: 4
Number of external LSAs 0, checksum sum 0
Number of areas is: 2, 2 normal, 0 stub
Number of active areas is: 2, 2 normal, 0 stub
Area (0.0.0.0) (Active)
  Interfaces in this Area: 1 Active Interfaces: 1
  Passive Interfaces: 0 Loopback Interfaces: 0
  SPF calculation has run 5 times
  Number of LSAs: 5, checksum sum 138261
Area (0.0.0.1) (Active)
  Interfaces in this Area: 1 Active Interfaces: 1
  Passive Interfaces: 0 Loopback Interfaces: 0
  SPF calculation has run 5 times
  Area ranges:
    fd00::1/64, inter-area, no-advertise
  Number of LSAs: 4, checksum sum 85028

Routing Process 2 with ID: 1.1.1.7 VRF red
------------------------------------------

Graceful-restart is configured
Restart Interval: 120, State: inactive
Last Graceful Restart Exit Status: none
Maximum Paths to Destination: 4
Number of external LSAs 0, checksum sum 0
Number of areas is: 1, 1 normal, 0 stub
Number of active areas is: 1, 1 normal, 0 stub
Area (0.0.0.1) (Active)
  Interfaces in this Area: 1 Active Interfaces: 1
  Passive Interfaces: 0 Loopback Interfaces: 0
  SPF calculation has run 4 times
  Number of LSAs: 4, checksum sum 149947


*********************************
Command : show ospfv3-af ipv6 interface all-vrfs
*********************************
Address-family : IPv6 Unicast
------------------------------

Interface 1/1/1 is up, Line Protocol is up
```

```
-----------------------------------------

IPv6 address fe80::7272:cfff:fea7:c8fa
Process ID 1 VRF default, Area 0.0.0.0
    State Backup-dr, Status up, Network type Broadcast, cost 1
    Transit delay 1 sec, Router priority 1
    Designated Router ID: 3.3.3.3
    Backup Designated Router ID: 1.1.1.3
    Timer Intervals: Hello 10, Dead 40, Retransmit 5
    Number of Link LSAs: 2, checksum sum 60947

Interface 1/1/2 is up, Line Protocol is up
-----------------------------------------

IPv6 address fe80::7272:cfff:fea7:c8fa
Process ID 2 VRF red, Area 0.0.0.1
    State Backup-dr, Status up, Network type Broadcast, cost 1
    Transit delay 1 sec, Router priority 1
    Designated Router ID: 4.4.4.4
    Backup Designated Router ID: 1.1.1.7
    Timer Intervals: Hello 10, Dead 40, Retransmit 5
    Number of Link LSAs: 2, checksum sum 75181


********************************
Command : show ospfv3-af ipv6 statistics interface all-vrfs
********************************
Address-family : IPv6 Unicast
------------------------------

OSPFv3 AF Process ID 1 VRF default, Interface 1/1/1 Statistics (cleared 0h 4m 45s
ago)

================================================================================
==

Tx Hello packets      : 28               Rx Hello packets      : 28
Tx Hello bytes        : 1116             Rx Hello bytes        : 1116
Tx DD packets         : 3                Rx DD packets         : 2
Tx DD bytes           : 164              Rx DD bytes           : 136
Tx LS request packets : 1                Rx LS request packets : 1
Tx LS request bytes   : 64               Rx LS request bytes   : 64
Tx LS update packets  : 4                Rx LS update packets  : 4
Tx LS update bytes    : 368              Rx LS update bytes    : 452
Tx LS ack packets     : 3                Rx LS ack packets     : 2
Tx LS ack bytes       : 208              Rx LS ack bytes       : 152

Total IPsec packets processed : 10
Total IPsec bytes processed   : 6400
Total Number of State Changes : 3
Number of LSAs                : 2
LSA Checksum Sum              : 60947

Total OSPFv3-AF Packets Discarded: 0
---------------------------------

Reason                          Packets Dropped
-----------------------------------------------
Invalid Type                    0
Invalid Length                  0
Invalid Version                 0
Bad or Unknown Source           0
```

```
Area Mismatch                    0
Self-originated                  0
Duplicate Router ID              0
Interface Standby                0
Total Hello Packets Dropped      0
  Hello Interval Mismatch        0
  Dead Interval Mismatch         0
  Options Mismatch               0
  MTU Mismatch                   0
  Neighbor Ignored               0
Resource Failures                0
Bad LSA Length                   0
Others                           0
IPsec Authentication Errors      0

Total LSAs Ignored : 0
-----------------------
Bad Type           : 0
Bad Length         : 0
Invalid Data       : 0
Invalid Checksum   : 0

OSPFv3 AF Process ID 2 VRF red, Interface 1/1/2 Statistics (cleared 0h 4m 45s ago)
================================================================================

Tx Hello packets      : 28               Rx Hello packets      : 28
Tx Hello bytes        : 1116             Rx Hello bytes        : 1116
Tx DD packets         : 3                Rx DD packets         : 2
Tx DD bytes           : 144              Rx DD bytes           : 136
Tx LS request packets : 1                Rx LS request packets : 1
Tx LS request bytes   : 64               Rx LS request bytes   : 52
Tx LS update packets  : 4                Rx LS update packets  : 4
Tx LS update bytes    : 332              Rx LS update bytes    : 452
Tx LS ack packets     : 3                Rx LS ack packets     : 2
Tx LS ack bytes       : 208              Rx LS ack bytes       : 132

Total IPsec packets processed : 10
Total IPsec bytes processed   : 6400
Total Number of State Changes : 3
Number of LSAs                : 2
LSA Checksum Sum              : 75181

Total OSPFv3-AF Packets Discarded: 0
-----------------------------------

Reason                           Packets Dropped
-----------------------------------------------
Invalid Type                     0
Invalid Length                   0
Invalid Version                  0
Bad or Unknown Source            0
Area Mismatch                    0
Self-originated                  0
Duplicate Router ID              0
Interface Standby                0
Total Hello Packets Dropped      0
  Hello Interval Mismatch        0
  Dead Interval Mismatch         0
  Options Mismatch               0
  MTU Mismatch                   0
  Neighbor Ignored               0
Resource Failures                0
```

```
Bad LSA Length                   0
Others                           0
IPsec Authentication Errors      0

Total LSAs Ignored : 0
----------------------
Bad Type           : 0
Bad Length         : 0
Invalid Data       : 0
Invalid Checksum   : 0


********************************
Command : show ospfv3-af ipv6 virtual-links all-vrfs
********************************
Address-family : IPv6 Unicast
-----------------------------

Virtual link to router 40.40.40.40 is up
----------------------------------------

Process ID 21 VRF default, Transit area 0.0.0.1
Transit delay 1 sec
Timer Intervals: hello 10, dead 40, retransmit 5
No authentication
Number of Link LSAs: 0, checksum sum 0
4 state changes

=====================================================
[End] Feature ospfv3-af
=====================================================


=====================================================
Show Tech commands executed successfully
=====================================================
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

summary-address

```
summary-address  <PREFIX/LENGHT> [no-advertise | tag <TAG-VALUE>]
no summary-address <PREFIX/LENGHT> [no-advertise | tag <TAG-VALUE>]
```

Description

Summarizes the external routes with the matching address and mask. When advertising this route, its metric is set to the lowest cost path from among the routes that were summarized.

The **no** form of this command disables route summarization.

This command only works for an ASBR (Autonomous System Boundary Router).

Parameter	Description
<PREFIX/LENGHT>	Specifies the IP prefix/mask for advertising the aggregate route.
no-advertise	Do not advertise the aggregate route. Suppress routes that match the specified prefix/mask pair.
tag *<TAG-VALUE>*	Specify the tag for the aggregate route. The summary prefix will be advertised along with the tag value in External LSAs. Range: 0 to 4294967295

Examples

Setting OSPFv3 AF route summarization for IPv4:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-ipv4-uc)# summary-address 10.1.0.0/16
switch(config-ospfv3-ipv4-uc)# summary-address 10.1.0.0/16 no-advertise
switch(config-ospfv3-ipv4-uc)# summary-address 10.1.0.0/16 tag 10
```

Disabling route summarization for IPv4:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-ipv4-uc)# no summary-address 10.1.0.0/16
switch(config-ospfv3-ipv4-uc)# no summary-address 10.1.0.0/16 no-advertise
switch(config-ospfv3-ipv4-uc)# no summary-address 10.1.0.0/16 tag 10
```

Setting OSPFv3 AF route summarization for IPv6:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-ipv6-uc)# summary-address fd00::1/64
switch(config-ospfv3-ipv6-uc)# summary-address fd00::1/64 no-advertise
switch(config-ospfv3-ipv6-uc)# summary-address fd00::1/64 tag 10
```

Disabling route summarization for IPv6:

```
switch(config)# router ospfv3 1
switch(config-ospfv3-ipv6-uc)# no summary-address fd00::1/64
switch(config-ospfv3-ipv6-uc)# no summary-address fd00::1/64 no-advertise
switch(config-ospfv3-ipv6-uc)# no summary-address fd00::1/64 tag 10
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-ospfv3-ipv4-uc<PROCESS-ID>` `config-ospfv3-ipv6-uc<PROCESS-ID>`	Administrators or local user group members with execution rights for this command.

timers lsa-arrival

```
timers lsa-arrival <DELAY>
no timers lsa-arrival
```

Description

Configures the minimum delay between receiving the same LSA from a peer. The same LSA is an LSA that contains the same LSA ID number, LSA type, and advertising router ID. If an instance of the same LSA arrives before the delay expires, the LSA is dropped. Generally, the LSA arrival timer should be set to a value less than or equal to the start-time value for the command **timers throttle lsa** on the neighbor.

The **no** form of this command sets the LSA timers to default values.

Parameter	Description
<DELAY>	Specifies the delay in milliseconds. Range: 0 to 600000. Default: 1000.

Examples

Setting the LSA arrival timer:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# timers lsa-arrival 10
```

Setting the LSA arrival timer to default:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# no timers lsa-arrival
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-af-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

timers throttle lsa

```
timers throttle lsa start-time <START-TIME> hold-time <HOLD-TIME> max-wait-time <WAIT-TIME>
no timers throttle lsa
```

Description

Configures the timers for LSA generation.

The **no** form of this command sets the LSA timers to default values.

Parameter	Description
start-time *<START-TIME>*	Specifies the initial wait time in milliseconds after which LSAs are generated. When set to 0, the LSAs are generated without any delay. Range: 0 to 600000. Default: 5000.
hold-time *<HOLD-TIME>*	Specifies the amount of time, in milliseconds, between regeneration of an LSA. The hold time doubles each time the same LSA must be regenerated, until **max-wait-time** is reached. When set to 0, LSA regeneration time is not increased. Range: 0 to 600000. Default: 0.
max-wait-time *<WAIT-TIME>*	Specifies the maximum wait time, in milliseconds, for regeneration of the same LSA. When set to 0, LSA regeneration time is not increased. Range: 0 to 600000. Default: 0.

Examples

Setting the LSA timers:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# timers throttle lsa start-time 100 hold-time 1000 max-
wait-time 10000
```

Setting LSA timers to default values:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# no timers throttle lsa
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-ospfv3-af-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

timers throttle spf

```
timers throttle spf start-time <milliseconds> hold-time <milliseconds>
max-wait-time <milliseconds>
no timers throttle spf
```

Description

Configures timers for SPF calculation. There are three timers:

- **start-time**: Is the initial delay before an SPF calculation is started. Default is 200 milliseconds.
- **hold-time**: Is the progressive backoff time to wait before next scheduled SPF calculation. Default is 1000 milliseconds. If a route change event occurs during this period, the value doubles until it reaches the *max-wait-time*.
- **max-wait-time:** Is the maximum time to wait before the next scheduled SPF calculation. Default is 5000 milliseconds. This is used to limit the SPF hold timer and also defines the time to be considered for which the OSPF LSDB has to be stable, after which the SPF throttle mechanism is reset.

The **no** form of this command sets all the configured non-default timers to default value.

Parameter	Description
<START-TIME>	Time in milliseconds to set timer for initial SPF delay. Default: 200. Range: 1-600000 *Required*
<HOLD-TINME>	Time in milliseconds to set the minimum hold time between two consecutive SPF calculations. Default: 1000. Range: 1-600000 *Required*
<WAIT-TIME>	Time in milliseconds to set the maximum wait time between two consecutive SPF calculations. Default: 5000. Range: 1-600000 *Required*

Examples

Setting non-default timer values for SPF throttling:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# timers throttle spf start-time 100 hold-time 1000 max-
wait-time 10000
```

Setting default timer values for SPF throttling after configuring non-default values:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# no throttle spf
switch(config-ospfv3-af-1)# no timers throttle spf start-time 100 hold-time 1000
max-wait-time 10000
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100 8325 8325H 8325P 8360 9300	config-ospfv3-af-*<PROCESS-ID>*	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
9300S 10000		

transit-delay

```
transit-delay <seconds>
no transit-delay
```

Description

Sets the time delay in Link state transmission for virtual links.

The **no** form of this command sets the delay in Link state transmission to the default of 1 second for virtual links.

Parameter	Description
<seconds>	Specifies the time delay for the transit delay, in seconds. Range: 1 to 3600. Default: 1.

Examples

Setting OSPFv3 AF virtual links transit delay:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 100 virtual-link 100.0.1.1
switch(config-router-vlink6)# transit-delay 30
```

Setting OSPFv3 AF virtual links transit delay to default:

```
switch(config)# router ospfv3-af 1
switch(config-ospfv3-af-1)# address-family ipv6 unicast
switch(config-ospfv3-ipv6-uc)# area 100 virtual-link  100.0.1.1
switch(config-router-vlink6)# no transit-delay
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L774 6400 8100	config-router-vlink	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325 8325H 8325P 8360 9300 9300S 10000		

Chapter 7
Border Gateway Protocol (BGP)

BGPv4 (RFC 4271) is an interdomain path-vector routing protocol that runs over TCP on port 179. It provides a reliable loop-free routing between different Autonomous Systems (ASs). BGP routers (called speakers) form peerings (also known as neighbor relationships) with other BGP peers. These peerings can either be internal (iBGP: within the same AS) or external, (eBGP: connecting two different ASs).

Overview

The characteristics of BGP are:

- Controls route propagation and the selection of optimal routes, rather than route discovery and calculation. This makes BGP different from interior gateway protocols (IGP) such as OSPF and RIP.
- Uses TCP to enhance reliability.
- Supports CIDR.
- Reduces bandwidth consumption by advertising only incremental updates, which allows advertising large amounts of routing information on the Internet.
- Eliminates routing loops by adding AS path information to BGP routes.
- Provides policies to implement flexible route filtering and selection.
- Provides scalability.

BGP views an autonomous system as a collection of routes under the control of a single organization, using one or more IGPs to route packets within the AS. The IGP could also be an iBGP within the AS and could use BGP as the only routing protocol.

Autonomous system numbers

Organizations requiring connectivity to the Internet must obtain an Autonomous System Number (ASN). ASNs were originally 2 bytes (16 bit), providing 65,535 ASNs. ASNs 64,512-65,534 are private ASNs within the 16-bit ASN range.

BGP sessions

A BGP session refers to the established adjacency between two BGP routers. BGP sessions are always point-to-point and are categorized into two types:

- **Internal BGP (IBGP)**: Sessions established with an IBGP router that are in the same AS. IBGP prefixes are assigned an administrative distance (AD) of 200 upon installing into the router routing information base (RIB).
- **External BGP (EBGP)**: Sessions established with a BGP router that are in a different AS. EBGP prefixes are assigned an AD of 20 upon installing into the router RIB.

Inter-router communication

BGP does not use hello packets to discover neighbors like IGP protocols and cannot discover neighbors dynamically. BGP was designed as an inter-autonomous routing protocol, implying neighbor adjacencies will not change frequently and are manually configured. BGP neighbors are defined by an IP address.

BGP uses TCP port 179 to communicate with other routers. While BGP can form neighbor adjacencies that are directly connected, it can also form adjacencies that are multiple hops away. Multihop sessions require that the router use an underlying route installed in RIB to establish the TCP session with a remote endpoint.

In the following scenario, R1 is able to establish a direct BGP session with R2. In addition, R2 is able to form a BGP session with R4, even though it passes through R3.

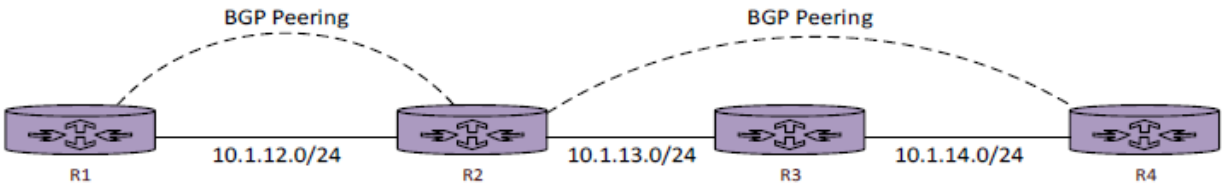

BGP messages

BGP communication uses four message types, as shown in the following table:

Type	Name	Functional Overview
1	OPEN	Sets up and establishes BGP adjacency
2	UPDATE	Advertises, updates, or withdraw routes
3	NOTIFICATION	Indicates and error condition to a BGP neighbor
4	KEEPALIVE	Ensures that BGP neighbors are still alive

BGP neighbor states

BGP forms a TCP session with neighbor routers called peers. The BGP sessions may report in the following states:

State	Listen for TCP	Initiate TCP	TCP Up	Open Sent	Open Received	Neighbor Up
Idle	No					

State	Listen for TCP	Initiate TCP	TCP Up	Open Sent	Open Received	Neighbor Up
Connect	Yes					
Active	Yes	Yes				
Open sent	Yes	Yes	Yes	Yes		
Open confirm	Yes	Yes	Yes	Yes	Yes	
Established	Yes	Yes	Yes	Yes	Yes	Yes

Injecting routes/prefixes into the BGP table

The BGP Routing Information Base (RIB) holds the network layer reachability information (NLRI) learned by BGP, as well as the associated path attributes (PAs). An NLRI is simply an IP prefix and prefix length.

A BGP router adds entries to its local BGP table by using the `network` command. The router receives the prefixes through an Update message from a neighbor or by redistributing from another routing protocol.

The network command instructs BGP protocol to perform the following steps:

- Look for a route in the current IP routing table that matches the parameters of the `network` command. If the IP route exists, put the equivalent NLRI into the local BGP table.
- With this logic, connected routes, static routes, or IGP routes can be taken from the IP routing table and placed into the BGP table for later advertisement.

The BGP `redistribute` command can redistribute static, connected, and IGP learned routes.

Path attributes

BGP attaches path attributes (PAs) associated with each network path. The PAs provide BGP with granularity and control of routing policies within BGP. The BGP prefix PAs are classified as follows:

- Well-known mandatory
- Well-known discretionary
- Optional transitive
- Optional non-transitive

Per RFC 4271, well-known attributes must be recognized by all BGP implementations. Well-known mandatory attributes must be included with every prefix advertisement, while well-known discretionary attributes may or may not be included.

Optional attributes do not have to be recognized by all BGP implementations. Optional attributes can be set to be transitive and stay with the route advertisement from AS to AS. Other PAs are nontransitive and cannot be shared from AS to AS. BGP PAs are summarized in the following table.

Path Attribute	Description	Characteristics
AS_PATH	Lists ASNs through which the route has been advertised.	Well-known mandatory

Path Attribute	Description	Characteristics
NEXT_HOP	Lists the next-hop IP address used to reach an NLRI.	Well-known mandatory
LOCAL_PREF	Used to communicate a BGP router degree of preference for an advertised route. Used only in updates between internal BGP peers.	Well-known discretionary
MULTI_EXIT_DISc	Known as the MED, used to influence incoming traffic only in EBGP updates. Allows an AS to inform another AS of its preferred ingress point.	Optional non-transitive
AGGREGATOR	Lists the RouterId and ASN of the router that created the summary NLRI and indicates where the Aggregation occurred.	Optional transitive
ATOMIC_AGGREGATOR	Tags a summary NLRI as being a summary and indicates that a loss of path information has occurred.	Well-known discretionary
ORIGIN	Describes where the route was imported into BGP; i (IGP), e (EGP), or ? (incomplete information).	Well-known mandatory
COMMUNITY	Identifies a destination as a member of some community of destinations that share one or more common properties. Used to simplify policy enforcement.	Optional transitive
ORIGINATOR_ID	Used by route reflectors (RRs) to denote the router-id of the iBGP neighbor that injected the NLRI into the AS	Optional non-transitive
CLUSTER_LIST	Used by route reflectors (RRs) to list a sequence of route reflection cluster IDs through which the route has passed in order to prevent loops.	Optional non-transitive
MP_REACH_NLRI	Carries the set of reachable destinations together with the next hop information to be used for forwarding to these destinations.	Optional non-transitive
MP_UNREACH_NLRI	Used to withdraw unreachable routes.	Optional non-transitive

BGP best-path calculation

When a BGP router learns multiple routes to the same NLRI, it must choose a single best route. The following list defines the BGP route selection decision process.

1. **Highest administrative weight**: Administrative weight can be assigned to each NLRI locally on a router and the value cannot be communicated to another router. A higher value indicates a better route.
2. **Highest LOCAL_PREF PA**: This attribute can be set on a router inside an AS and distributed inside the AS. A higher value indicates a better route.
3. **Locally injected routed**: Pick the route injected into BGP locally using the network command `redistribution`.

4. **Shortest AS_PATH length**: A shorter AS_PATH indicates a better route. This attribute counts each ASN in the AS_SEQUENCE as one.
5. **Origin PA**: IGP (I) routes are preferred over EGP (E) routes, which are in turn preferred over incomplete (?) routes.
6. **Smallest Multi-Exit Discriminator (MED) PA**: A smaller value indicates a better route.
7. **Neighbor Type**: External BGP (eBGP) routes are better than internal BGP (iBGP) routes.
8. **IGP metric for reaching the NEXT_HOP**: When comparing IGP metrics for each NLRIs NEXT_HOP, a smaller value indicates a better route.
9. **Keep the oldest eBGP route**: If the routes being compared are eBGP and one is the best path, retain the existing best path. This action reduces eBGP route flaps.
10. **Choose the smallest neighbor router ID**: Use the route with the smallest next-hop router router-id. Only perform this step if bgp bestpath compare-routerid is configured.
11. **Smallest neighbor ID**: The local router has at least two neighbor relationships with a single other router. For this case, the router prefers the route advertised by the lowest neighbor ID.

Loop prevention

BGP is a path vector routing protocol and does not contain a complete topology of the link state routing protocols.

The BGP attribute AS_PATH is a mandatory attribute and includes a complete listing of all ASNs that the prefix advertisement has traversed from its source AS. AS_PATH is used as a loop prevention mechanism in the BGP protocol. If a BGP router receives a prefix advertisement with its AS listed in the AS_PATH, it discards the prefix because the router thinks the advertisement forms a loop.

Route policies

Route policy is a method that defines how routes are accepted on a router and how routes are advertised. Filtering of routes within BGP is accomplished with route maps, prefix lists, aspath lists, or community lists.

Resetting BGP sessions

When a BGP router configuration is changed, it is often necessary to reset the connections to the affected neighbors in order for the change to take effect. The user must reset the BGP connection for all:

- Changes to BGP related route maps, prefix lists, aspath lists, and community-lists
- Changes to the BGP-related timer specification
- Changes to BGP-related weights

BGP supports two methods of clearing a BGP session:

- A hard reset, which tears down the BGP session, removes BGP routes from the peer, and is the most disruptive.
- A soft reset, which invalidates the BGP cache and requests a full advertisement from its BGP peer.

Initiate a hard or soft reset with the `clear bgp` commands. Clear sessions with all BGP neighbors by using an asterisk `*` instead of the peer IP address.

When a BGP route policy changes, the BGP table must be processed again to notify neighbors. If the BGP session supports route refresh capability, the peer refreshes the prefixes to the requesting router,

allowing the inbound policy to process using the new policy changes. The route refresh capability is negotiated when the session is established.

If the BGP session does not support route refresh, use inbound soft reconfiguration so that updates to inbound route policies can be applied without performing a hard reset.

- Inbound soft reconfiguration does not purge the Adj- RIB-In table after routes process into the Loc-RIB table.
- The Adj-RIB-In table maintains the raw unedited routes received from neighbors and allows the inbound route policies to be processed again.
- Enabling this feature can consume a significant amount of memory because the Adj-RIB-In table stays in memory.

IBGP full mesh requirement

BGP uses the AS_PATH as a loop detection and prevention mechanism because the ASN is prepended when advertising to an EBGP neighbor. IBGP peers do not prepend their ASN to the AS_PATH because the NLRIs would fail the validity check and would not install the prefix in the IP routing table.

No other method exists to detect loops with IBGP sessions. RFC 4271 prohibits the advertisement of an NLRI received from one IBGP peer to another. RFC 4271 states that all BGP routers within a single AS must be fully meshed to provide a complete loop-free routing table and prevent traffic null holing.

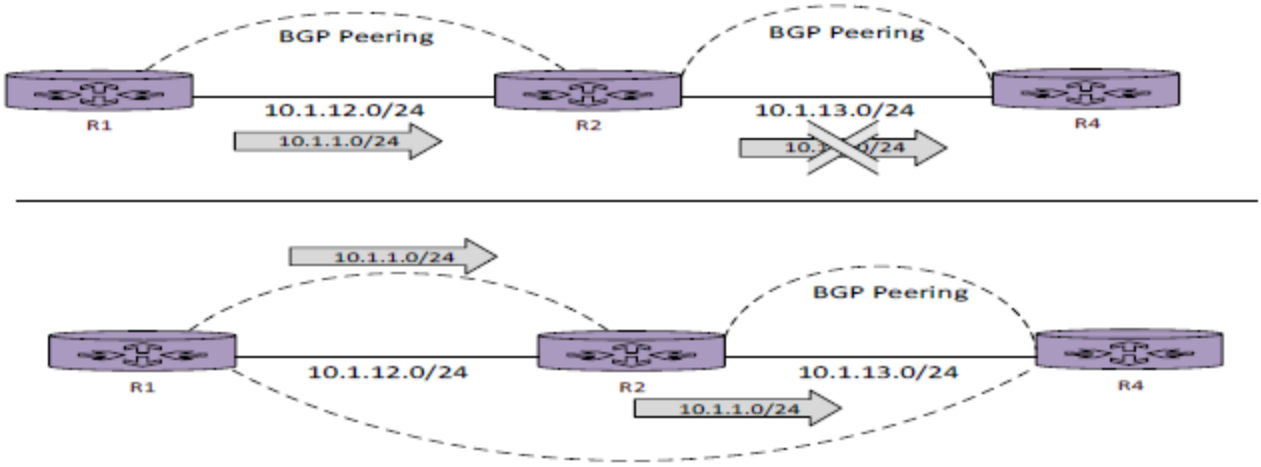

In the figure above, R1, R2, and R3 are all within AS65100.

- R1 has an IBGP session with R2 and R2 has an IBGP session with R3.
- R1 advertises the 10.1.1.0/24 prefix to R2, which is processed and inserted into the R2 BGP table.
- R2 does not advertise the 10.1.1.0/24 NLRI to R3 because it received the prefix from an IBGP peer.
- To resolve this issue, R1 must form a multi-hop IBGP session so that R3 can receive the 10.1.1.0/24 prefix directly from R1.

Route reflectors

The inability of BGP to advertise a prefix learned from one IBGP peer to another can lead to scalability issues within an AS. The formula n(n-1)/2 provides the number of sessions required where n represents the number of routers. A full mesh topology of 10 routers requires 45 sessions. IBGP scalability becomes an issue for large networks.

RFC 1966 specifies that IBGP peering can be configured so that it reflects routes to another IBGP peer. The router reflecting route is known as a route reflector (RR) and the router receiving reflected routes is a route reflector client. Three basic rules apply to route reflectors and route reflection

1. If an RR receives an NLRI from a non-RR client, the RR advertises the NLRI to an RR client. It does not advertise the NLRI to a nonroute reflector client.
2. If an RR receives an NLRI from an RR client, it advertises the NLRI to RR and non-RR clients. Even the RR client that sent the advertisement receives a copy of the route, but it discards the NLRI.
3. If an RR receives a route from an EBGP peer, it advertises the route to RR and non-RR clients.

Loop prevention in route reflectors

Removing the full mesh requirement in an IBGP topology introduces the potential for routing loops. RFC 1966 introduces two other BGP route reflector-specific attributes to prevent loops.

ORIGINATOR_ID is an optional non-transitive BGP attribute created by the first route reflector. This attribute sets the value to the router-id of the router that injected/advertised the route into the AS. If a router receives an NLRI with its router-id in the Originator attribute, the NLRI is discarded.

CLUSTER_LIST is a non-transitive BGP attribute updated by the route reflector. This attribute is appended by the route reflector with its cluster-id. By default this ID is the BGP router-id. The cluster-id can be set with BGP configuration command `bgp cluster-id` **`<cluster-id>`**. If a router receives an NLRI with its cluster-id in the cluster list attribute, the NLRI is discarded.

BGP unnumbered

BGP unnumbered is designed to ease and automate the configuration in the datacenter by eliminating the need for interface IP address and specifying the **remote-as** of each neighbor. BGP unnumbered is configured by specifying the interface name to which the peer is directly connected. BGP, however, needs an IP address to create TCP/IP connection. BGP makes use of the neighbor's IPv6 link-local, which is learned via the router advertisement.

If the BGP unnumbered is configured over an SVI interface, then only one interface is expected to have the router connected to it.

With SVI, BGP peering is only established with one peer at a time.

BGP peer groups

Peer groups can be used to assign common policies and attributes such as an AS number or source-interface for multiple neighbors. A peer group is relevant only to the router on which it is active and is not communicated to router peers. A BGP neighbor can belong to only one peer group. Peer-group activate is not supported in the current release. Peers under the peer groups must be activated individually under the address families.

BGP communities

BGP communities provide additional capability for tagging routes and modifying BGP routing policy on upstream and downstream routers. They can be appended, removed, or modified on each attribute as the route travels from router to router.

BGP communities are an optional transitive BGP attribute that can traverse from autonomous system to autonomous system. A BGP community is a 32-bit number that can be included with a route. It can be displayed as two 16-bit numbers (0-65535):(0-65535).

Aggregate routes

Aggregating routes conserves router resources and accelerates best patch calculation by reducing the size of the table. The two methods for BGP route aggregation are:

- **Static**: Create a static reject route for the prefix and advertise the network through a `network` command. The disadvantage of this method is that the summary route will always be advertised even if the networks are not available.
- **Dynamic**: Configure an aggregation network range using the `aggregate-address` command. When viable routes that match the network range enter the BGP table, an aggregate route is created. On the originating router, a reject route for the aggregated prefix is automatically created by BGP as a loop-prevention mechanism.

BGP Graceful-Restart and high availability

The BGP Graceful-Restart (GR) feature allows a BGP speaker to express its ability to preserve forwarding state during:

- BGP protocol (or daemon) restart
- Management Module (MM) switchover

BGP GR is enabled by default and announces GR capability in the BGP OPEN message to peers. BGP initiates the graceful-restart process when an MM switchover occurs and also acts as a GR-aware device.

A GR-aware device, also known as GR helper mode, is notified that the peer router is transitioning and takes appropriate actions based on configuration or default timers.

When a BGP restart happens on the peer router or when MM switchover occurs, the routes currently held in the forwarding table are marked as stale. Thus the forwarding state is preserved as the control plane and the forwarding plane operates independently.

On the restarting peer on which the switchover occurred, BGP on the newly active MM starts to establish sessions with all configured peers. BGP on the nonrestarting side sees new connection requests coming in while BGP is in an established state. Such an event is an indication for the nonrestarting peer that the peer has restarted. At this point, the restarting peer sends the GR capability with Restart bit set to 1 and Forwarding State bit set to 1.

The nonrestarting peer:

- Cleans up old BGP sessions and marks all the routes in the BGP table that are received from the restarting peer as stale
- Purges all stale routes after the Restart Time expires (if the restarting peer never re-establishes the BGP session)
- Sends an initial routing table update, followed by an End-of-RIB (EoR)

The restarting peer:

- Delays best-path calculation until after receiving EoR from all peers
- Generates updates for its peers and sends the EoR marker after the initial table is sent

The nonrestarting peers:

- Receive the routing updates from the restarting peer
- Remove stale marking for any refreshed route
- Purge any remaining stale routes after EoR is received from the restarting peer or the stale path timer expires

Configuration tasks list

Configuring BGP peers

Assign routes to the AS numbers using the router bgp command. For example, the following configuration assigns the router RT1 to AS 100:

```
switch(config)# router bgp 100
```

This other configuration assigns the router RT2 to AS 200 :

```
switch(config)# router bgp 200
```

Defining BGP Neighbors

After assigning each router to an AS, define BGP neighbors using the following commands:

- neighbor remote-as
- address-family
- neighbor activate

Sample configuration:

```
switch(config)# router bgp 100
switch(config-bgp)# neighbor 192.168.1.2 remote-as 200
switch(config-bgp)# address-family ipv4 unicast
switch(config-bgp-ipv4-uc)# neighbor 192.168.1.2 activate
```

Using Loopback Interfaces for BGP peering

If the IP address used in the neighbor command is not directly attached to the interface, then the `update-source`command must be used to specify the source IP address for the TCP connection. This source IP address can be a loopback, IP address, LAG, VLAN, or other interface that is directly connected to the router.

Sample configuration:

```
router bgp <AS>
neighbor <IP-ADDRESS> update-source
 A.B.C.D IPv4 address
 X:X::X:X IPv6 address
 lag Lag Interface
 loopback Loopback Interface
 vlan Vlan Interface
```

Verifying BGP routing table

To verify the BGP routing table, use the following commands:

For default VRF: `show bgp {ipv4 unicast | ipv6 unicast | all } [vsx-peer]`

For non-default VRF: `show bgp [vrf <VRF-NAME>] {ipv4 unicast | ipv6 unicast | all } [vsx-peer]`

Sample output:

```
switch#  show bgp vrf v1 ipv4 unicast

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete
VRF : v1
Local Router-ID 9.0.0.1
Network Nexthop Metric LocPrf Weight Path
------------------------------------------------------------------------
*>e 9.0.0.0/24 9.0.0.2 0 100 0 65534.65535 3.4
18.54934 3574.8570 5.6 ?
*>e 100.0.0.0/24 9.0.0.2 0 100 0 200 ?
*>e 100.0.1.0/24 9.0.0.2 0 100 0 200 ?
*>e 100.0.2.0/24 9.0.0.2 0 100 0 200 ?
*>e 100.0.3.0/24 9.0.0.2 0 100 0 200 ?
*ae 100.0.3.0/24 9.0.0.3 0 100 0 200 ?
Total number of entries 6
```

Verifying BGP peers

To verify the BGP peers, use the following commands:

For default VRF: `show bgp {ipv4 unicast | ipv6 unicast | all} neighbors [vsx-peer]`

For non-default VRF: `show bgp [vrf <VRF-NAME>] {ipv4 unicast | ipv6 unicast | all} neighbors [vsx-peer]`

Sample output:

```
switch# show bgp all neighbors
Codes: ^ Inherited from peer-group, * Dynamic Neighbor
VRF : default

BGP Neighbor 10.1.1.2 (Internal)
    Description         :
    Peer-group          :

    Remote Router Id    : 10.1.1.2          Local Router Id    : 1.0.0.1
    Remote AS           : 1                 Local AS           : 1
    Remote Port         : 0                 Local Port         : 0
    State               : Idle              Admin Status       : Up
    Conn. Established   : 0                 Conn. Dropped      : 0
    Passive             : No                Update-Source      :
    Cfg. Hold Time      : 180               Cfg. Keep Alive    : 60
    Neg. Hold Time      : 0                 Neg. Keep Alive    : 0
    Up/Down Time        : 00h:00m:00s       Alt. Local-AS      : 0
    Local-AS Prepend    : No
    BFD                 : Disabled
    Ignore-Leading-AS   : No
    TCP AO Keychain     : kc1
    TCP AO Current Key  : 10
    Send-ID             : 218               Recv-ID            : 218
    Include TCP Options : yes               Accept AO Mismatch : yes
    Last Err Sent       : No Error
    Last SubErr Sent    : No Error
    Last Err Rcvd       : No Error
    Last SubErr Rcvd    : No Error

    Graceful-Restart    : Enabled           Gr. Restart Time   : 120
    Gr. Stalepath Time  : 150               Remove Private-AS  : No
    TTL                 : 255               Local Cluster-ID   :
    Weight              : 0                 Fall-over          : No
```

```
  Message statistics       Sent     Rcvd
  -------------------      -----    -----
  Open                        0        0
  Notification                0        0
  Updates                     0        0
  Keepalives                  0        0
  Route Refresh               0        0
  Total                       0        0

  Capability                          Advertised      Received
  ----------                          ----------      ----------
  Route Refresh                       Yes             No
  Graceful Restart                    Yes             No
  Four Octet ASN                      Yes             No
  Address family IPv4 Unicast         Yes             No
  Address family IPv6 Unicast         No              No
Address Family : IPv4 Unicast
  ----------------------------

  Rt. Reflect. Client : No                 Send Community    :
  Allow-AS in         : 0                  Advt. Interval    : 30
  Max. Prefix         : 64000              Soft Reconfig In  :
  Nexthop-Self        :                    Default-Originate :

  Routemap In         :
  Routemap Out        :
  Address Family : IPv6 Unicast
  ----------------------------
```

Clearing BGP neighbors

After establishing a BGP peer, if you need to make changes to the configuration of the BGP neighbor sub-command such as next-hop-self, route-map, or send-community, you must reset neighbor connection for new configuration to take into effect,

The following commands clears the specified neighbors:

```
Default VRF: clear bgp <NEIGHBOR-IP-ADDRESS>
             clear bgp ipv4 unicast 192.168.12.1 soft in
Non-default VRF: clear bgp vrf <VRF-NAME> <NEIGHBOR-IP-ADDRESS>
```

The following commands clear all neighbors in the default VRF and the non-default VRF, respectively:

```
Defualt VRF: clear bgp *
Non-Default VRF: clear bgp vrf <VRF-NAME> a *
```

Address families

MP-BGP adds support for Address families and Sub Address families in addition to IPv4 Unicast which was supported by base BGP (RFC 4271). The 10.02 release supports IPv4 Unicast and IPv6 Unicast address families. Use the `address-family {<ipv4> | <ipv6>} unicast` commands to enter the address-family configuration mode and execute address-family specific configurations.

Scale limits

- Maximum neighbors supported across all VRFs: 256
- Maximum equal cost paths supported: 8
- Maximum route reflector clients allowed across all VRFs: 256
- Maximum BGP routes supported: 256,000
- Maximum routes accepted from a BGP peer: 256,000
- Maximum peer groups allowed across all VRFs: 50
- Maximum AS numbers in as-path attribute: 32
- Maximum as-path entries in a single aspath-list: 128
- Maximum aspath-lists: 128
- Maximum community entries in a single community-list: 128
- Maximum community-lists: 256
- Maximum prefix entries in a single prefix-list: 128
- Maximum prefix-lists: 256
- Maximum route map entries in a single route-map: 128
- Maximum route-maps: 256

Configuration tasks list

Configuring BGP peers

Assign routes to the AS numbers using the router bgp command. For example, the following configuration assigns the router RT1 to AS 100:

```
switch(config)# router bgp 100
```

This other configuration assigns the router RT2 to AS 200 :

```
switch(config)# router bgp 200
```

Defining BGP Neighbors

After assigning each router to an AS, define BGP neighbors using the following commands:

- neighbor remote-as
- address-family
- neighbor activate

Sample configuration:

```
switch(config)# router bgp 100
switch(config-bgp)# neighbor 192.168.1.2 remote-as 200
switch(config-bgp)# address-family ipv4 unicast
switch(config-bgp-ipv4-uc)# neighbor 192.168.1.2 activate
```

Using Loopback Interfaces for BGP peering

If the IP address used in the neighbor command is not directly attached to the interface, then the `update-source`command must be used to specify the source IP address for the TCP connection. This source IP address can be a loopback, IP address, LAG, VLAN, or other interface that is directly connected to the router.

Sample configuration:

```
router bgp <AS>
neighbor <IP-ADDRESS> update-source
 A.B.C.D IPv4 address
 X:X::X:X IPv6 address
 lag Lag Interface
 loopback Loopback Interface
 vlan Vlan Interface
```

Verifying BGP routing table

To verify the BGP routing table, use the following commands:

For default VRF: `show bgp {ipv4 unicast | ipv6 unicast | all } [vsx-peer]`

For non-default VRF: `show bgp [vrf <VRF-NAME>] {ipv4 unicast | ipv6 unicast | all } [vsx-peer]`

Sample output:

```
switch#  show bgp vrf v1 ipv4 unicast

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete
VRF : v1
Local Router-ID 9.0.0.1
Network Nexthop Metric LocPrf Weight Path
------------------------------------------------------------------------
*>e 9.0.0.0/24 9.0.0.2 0 100 0 65534.65535 3.4
18.54934 3574.8570 5.6 ?
*>e 100.0.0.0/24 9.0.0.2 0 100 0 200 ?
*>e 100.0.1.0/24 9.0.0.2 0 100 0 200 ?
*>e 100.0.2.0/24 9.0.0.2 0 100 0 200 ?
*>e 100.0.3.0/24 9.0.0.2 0 100 0 200 ?
*ae 100.0.3.0/24 9.0.0.3 0 100 0 200 ?
Total number of entries 6
```

Verifying BGP peers

To verify the BGP peers, use the following commands:

For default VRF: `show bgp {ipv4 unicast | ipv6 unicast | all} neighbors [vsx-peer]`

For non-default VRF: `show bgp [vrf <VRF-NAME>] {ipv4 unicast | ipv6 unicast | all} neighbors [vsx-peer]`

Sample output:

```
switch# show bgp all neighbors
Codes: ^ Inherited from peer-group, * Dynamic Neighbor
VRF : default
```

```
BGP Neighbor 10.1.1.2 (Internal)
    Description          :
    Peer-group           :

    Remote Router Id     : 10.1.1.2          Local Router Id     : 1.0.0.1
    Remote AS            : 1                 Local AS            : 1
    Remote Port          : 0                 Local Port          : 0
    State                : Idle              Admin Status        : Up
    Conn. Established    : 0                 Conn. Dropped       : 0
    Passive              : No                Update-Source       :
    Cfg. Hold Time       : 180               Cfg. Keep Alive     : 60
    Neg. Hold Time       : 0                 Neg. Keep Alive     : 0
    Up/Down Time         : 00h:00m:00s       Alt. Local-AS       : 0
    Local-AS Prepend     : No
    BFD                  : Disabled
    Ignore-Leading-AS    : No
    TCP AO Keychain      : kc1
    TCP AO Current Key   : 10
    Send-ID              : 218               Recv-ID             : 218
    Include TCP Options  : yes               Accept AO Mismatch  : yes
    Last Err Sent        : No Error
    Last SubErr Sent     : No Error
    Last Err Rcvd        : No Error
    Last SubErr Rcvd     : No Error

    Graceful-Restart     : Enabled           Gr. Restart Time    : 120
    Gr. Stalepath Time   : 150               Remove Private-AS   : No
    TTL                  : 255               Local Cluster-ID    :
    Weight               : 0                 Fall-over           : No

    Message statistics       Sent     Rcvd
    ------------------       -----    -----
    Open                        0        0
    Notification                0        0
    Updates                     0        0
    Keepalives                  0        0
    Route Refresh               0        0
    Total                       0        0

    Capability                          Advertised      Received
    ----------                          ----------      ----------
    Route Refresh                       Yes             No
    Graceful Restart                    Yes             No
    Four Octet ASN                      Yes             No
    Address family IPv4 Unicast         Yes             No
    Address family IPv6 Unicast         No              No
   Address Family : IPv4 Unicast
    ----------------------------

    Rt. Reflect. Client : No                 Send Community    :
    Allow-AS in         : 0                  Advt. Interval    : 30
    Max. Prefix         : 64000              Soft Reconfig In  :
    Nexthop-Self        :                    Default-Originate :

    Routemap In         :
    Routemap Out        :
    Address Family : IPv6 Unicast
    ----------------------------
```

Clearing BGP neighbors

After establishing a BGP peer, if you need to make changes to the configuration of the BGP neighbor sub-command such as next-hop-self, route-map, or send-community, you must reset neighbor connection for new configuration to take into effect,

The following commands clears the specified neighbors:

```
Default VRF: clear bgp <NEIGHBOR-IP-ADDRESS>
             clear bgp ipv4 unicast 192.168.12.1 soft in
Non-default VRF: clear bgp vrf <VRF-NAME> <NEIGHBOR-IP-ADDRESS>
```

The following commands clear all neighbors in the default VRF and the non-default VRF, respectively:

```
Defualt VRF: clear bgp *
Non-Default VRF: clear bgp vrf <VRF-NAME> a *
```

Configuring BGP over an IPSec tunnel interface

BGP over an IPSec tunnel can be configured by defining the BGP peer as the IPSec tunnel overlay IP. The IP address of an IPSec interface is reachable over the IPSec tunnel.

This is supported only on the AOS-CX 10000 Switch Series.

BGP peering using IPSec tunnel interface IP

In the following example, BGP peering is configured between IPSec tunnel interface IPs, with **100.1.1.1** on **Switch 1** and **100.1.1.2** on **Switch 2**.

Switch 1:

```
interface tunnel 10 mode ipsec ipv4
    vrf attach vrf-green
    source ip 111.100.111.11
    destination ip 121.100.121.21
    ip address 100.1.1.1/24
    no shutdown


router bgp 1
!
    vrf vrf-green
        bgp router-id 100.1.1.1
        neighbor 100.1.1.2 remote-as 1
        address-family ipv4 unicast
            neighbor 100.1.1.2 activate
            network 60.60.60.0/24
        exit-address-family
!
```

Switch 2:

```
interface tunnel 10 mode ipsec ipv4
    vrf attach vrf-green
    source ip 121.100.121.21
```

```
    destination ip 111.100.111.11
    ip address 100.1.1.2/24
    no shutdown


router bgp 1
!
    vrf vrf-green
        bgp router-id 100.1.1.2
        neighbor 100.1.1.1 remote-as 1
        address-family ipv4 unicast
            neighbor 100.1.1.1 activate
            network 50.50.50.0/24
        exit-address-family
!
```

Verifying BGP peers:

On Switch 1:

```
Switch-1# show ip route vrf vrf-green

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF, D - DHCP
              U - Unnumbered
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2

VRF: vrf-green

Prefix              Nexthop        Interface     VRF       Origin/   Distance/
Age
                                                (egress)     Type      Metric
-------------------------------------------------------------------------------
-------------
50.50.50.0/24       100.1.1.2      tunnel10      -         B/I       [200/0]
00h:45m:17s
Total Route Count : 1
```

On Switch 2:

```
Switch-2# show ip route vrf vrf-green

Displaying ipv4 routes selected for forwarding

Origin Codes: C - connected, S - static, L - local
              R - RIP, B - BGP, O - OSPF, D - DHCP
              U - Unnumbered
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
              IA - OSPF internal area, E1 - OSPF external type 1
              E2 - OSPF external type 2

VRF: vrf-green

Prefix              Nexthop         Interface     VRF   Origin/   Distance/    Age
                                                (egress) Type      Metric
-------------------------------------------------------------------------------
```

```
---------------

60.60.60.0/24       100.1.1.1        tunnel10      -      B/I       [200/0]
00h:17m:05s


Total Route Count : 1
```

- It is recommended to use route-maps for BGP over IPsec sessions when there are BGP sessions providing connectivity between the tunnel endpoints in the underlay.
- You can have BGP sessions over IPsec tunnels using other sources, such as SVI or loopbacks. However, the reachability to these addresses should be through the static routes pointing to the IPsec tunnels.

Best practices and recommendations

The following are some best practices and recommendations to be followed when configuring BGP:

- On a 10000 Switch series, if a BGP session is formed over IPsec tunnels and the underlay routing protocol is also BGP, configure explicit route maps to prevent route leaking.
- It is recommended to configure iBGP using loopback interfaces, as the loopback interface is always up and reachable through the configured IP address.
- Configure router ID explicitly to ensure that a fixed router ID is used instead of relying on the highest IP address on loopback interfaces or physical interfaces.
- After applying inbound or outbound route maps to a BGP neighbor under the address-family context, a hard reset should be performed for the neighbor session to take effect.
- Configuring route summarization preserves router resources and reduces the size of the routing table.
- To advertise communities to other BGP peers, the `send-community` command should be enabled under the address-family.
- In a scenario where mutual redistribution is needed between IGP and BGP, it should be handled properly with route maps to make sure there is no routing loop.
- It is recommended to configure `next-hop-self` when setting up BGP sessions between VSX pairs. This will ensure that the next-hop for external routes is changed appropriately.

For more information, see the AOS-CX BGP over GRE Tunnels with IVRL video on the HPE Aruba Networking broadcasting channel on Youtube.

Use cases

The below design example shows the configuration and routing tables of AOS-CX routers:

Multi-homing to the same ISP

Scenario 1: Traffic from internal routers to reach ISP through Primary and Backup path (One link is primary and the other link is backup)

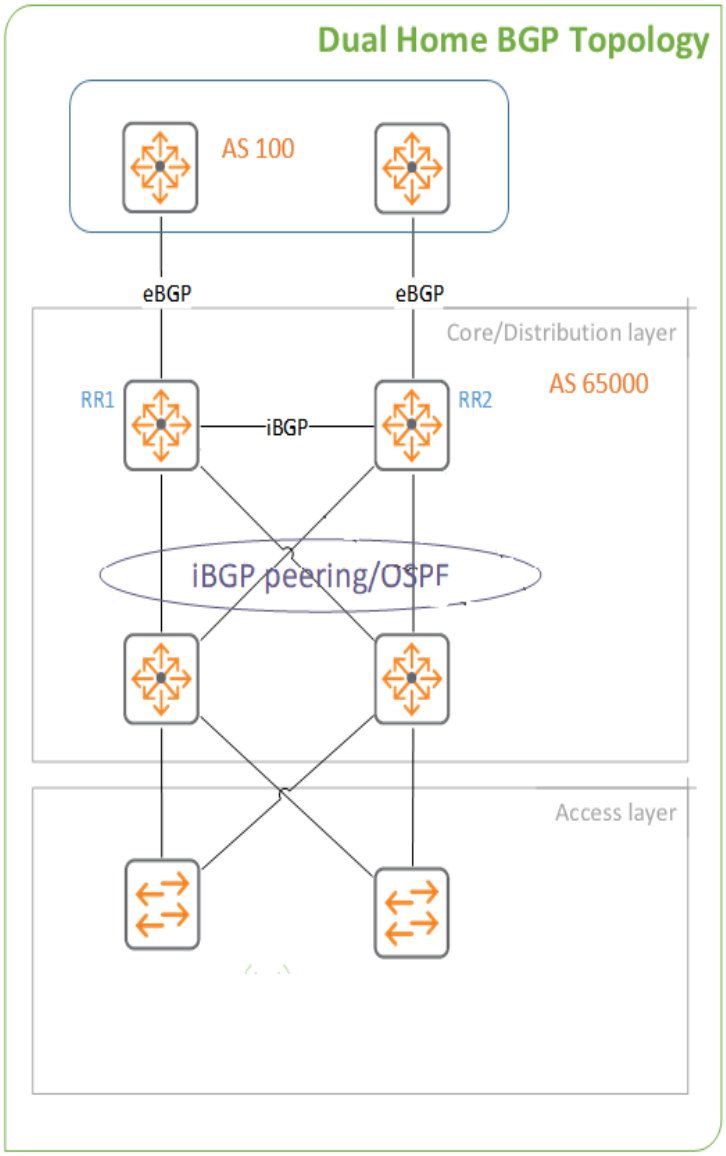

In the above topology, the RR1 and RR2 routers establish eBGP connections with two Provider Edge (PE) routers belonging to the same Internet Service Provider (ISP). In typical ISP scenarios, the ISP advertises a default route to the user, which allows all incoming and outgoing traffic to be routed through the primary link. In the event of a failure on the primary link, the secondary link can be used to ensure that traffic can still be routed to the Internet.

If the user runs BGP within their internal network, The RR1 and RR2 routers act as route-reflectors and are responsible for advertising the routes to reach outside organizations, such as the Internet, to the internal iBGP routers.

RR1 and RR2 routers should only advertise their internal network to the ISP.

BGP configuration:

RR1 Router	RR2 Router
`interface loopback 1`	`interface loopback 0`

RR1 Router

```
ip address 1.1.1.1/32
!
interface 1/1/1
ip address 192.168.11.1/24
!
router bgp 65000
bgp router-id 1.1.1.1
neighbor 192.168.11.2 remote-as 100
!
address-family ipv4 unicast
neighbor 192.168.11.2 activate
```

RR2 Router

```
ip address 2.2.2.2/32
!
interface 1/2/1
ip address 193.168.22.1/24
!
router bgp 65000
bgp router-id 2.2.2.2
neighbor 192.168.22.2 remote-as 100
!
address-family ipv4 unicast
neighbor 192.168.22.2 activate
```

RR2 router should also advertise all the internal prefixes with a longer AS_PATH list so that the PE routers will prefer the path with a shorter AS_PATH list. Hence, you must prepend to the ISP as given below:

Sample configuration on RR2:

```
route-map BGP-SET-AS permit seq 10
 set as-path prepend 300 47 50
!
router bgp 65000
 address-family ipv4 unicast
  neighbor 192.168.22.2 route-map BGP-SET-AS out

!
```

To ensure that internal routers in your network can reach the ISP and choose the exit path via RR1, while keeping the PE-RR2 link as a backup path, you can configure RR1 to advertise all routes received from the ISP PE router to all iBGP peers with a higher local-preference. This will make these routes more preferred than the routes advertised by the RR2 router. In the event of a failure or neighbor flap of RR1, traffic would then go through the RR2-PE link.

Sample configuration on RR1:

```
route-map BGP-SET-LP permit seq 10
 set local-preference 1000

!
router bgp 65000
 address-family ipv4 unicast
  neighbor <ip> route-map BGP-SET-LP out    << All internal iBGP peers
!
```

Scenario 2: Traffic from internal routers to reach ISP through both links (Load Balancing)

In most deployment modes, traffic from the internal network is expected to use both ISP links, as both are paid connections. By keeping the default configurations on RR1 and RR2 towards internal iBGP peers, you can achieve load-balancing towards external routes through both ISP links.

BGP AS-Path prepending and AS-Path filters

AS-Path prepending is a technique used in BGP to manipulate the AS-Path attribute of a BGP prefix. It involves adding multiple entries of AS to a BGP prefix. AS-Path prepending can be applied in both inbound and outbound directions using route-maps.

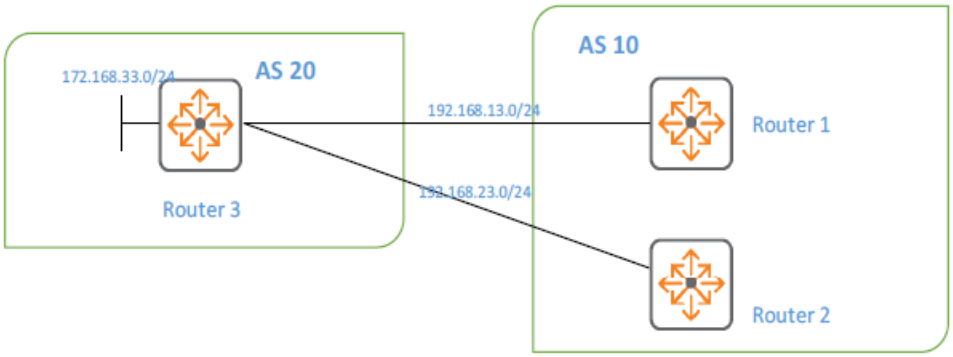

Outbound AS-Path prepending

AS-Path prepending can be applied to outbound direction on R3 router as shown below:

```
!
route-map test permit
set as-path prepend 20 20 20

router bgp 20
bgp router-id 1.1.1.1
neighbor 192.168.13.1 remote-as 10
neighbor 192.168.23.2 remote-as 10
address-family ipv4 unicast
neighbor 192.168.13.1 activate
neighbor 192.168.23.2 activate
neighbor 192.168.23.2 route-map test out
network 172.168.33.0/24
!
```

When R3 router advertises the 172.168.33.0/24 network to R2 router, it prepends its own AS 20 multiple times. However, R1 router receives the original BGP route with a single AS-path entry.

BGP tables on R1 and R2

```
R1# show bgp ipv4 unicast

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete
VRF : v1
Local Router-ID 9.0.0.1


    Network            Nexthop         Metric     LocPrf     Weight Path
-------------------------------------------------------------------------

*> 172.168.33.0/24      192.168.13.2      0          0          20    i
```

```
R2# show bgp ipv4 unicast

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete
VRF : v1
Local Router-ID 9.0.0.1


    Network            Nexthop        Metric     LocPrf     Weight Path
--------------------------------------------------------------------------
172.168.33.0/24       192.168.23.2     0            0         20 20 20 20 i
*>172.168.33.0/24     1.1.1.1 0        0            20          i
```

R2 router has two paths to 172.168.33.0/24. However, it chooses the path through R1 since that path has shortest AS-Path (indicated by a > sign, meaning best route). Hence, the direct path through R3 router can act as a backup path.

Inbound AS-Path Prepending

Inbound AS-Path prepending is configured using set as-path prepend last-as command under route-map which is then applied in inbound direction. The last-as keyword copies the AS number of the neighbor advertising the BGP route that matches the route-map.

Inbound AS-Path prepending on R1:

```
!
route-map test permit 10
set as-path prepend 20 20
!

router bgp 10
bgp router-id 2.2.2.2
neighbor 192.168.13.2 remote-as 20
address-family ipv4 unicast
neighbor 192.168.13.2 activate
neighbor 192.168.13.2 route-map test in
```

This configuration causes R1 to prepend the AS number of R3 twice before installing in the BGP table. R3 does not prepend any AS number by itself .

Debugging and troubleshooting

If BGP session is not established. use following debug flow to verify basic configurations:

1. Verify neighbor IP address configured under router BGP is reachable from DUT (using ping).
2. Verify BGP configurations on all DUTs.

```
router bgp 100
  bgp log-neighbor-changes
  bgp router-id 1.1.1.1
  neighbor 192.168.12.2 remote-as 100
  address-family ipv4 unicast
```

```
   neighbor 192.168 activate
exit-address-family
```

a. Manually set different BGP router-id in each DUT or ensure router-id is not same between peers. Verify BGP router-id using the `show bgp all summary` command.

```
switch# show bgp all summary

BGP Summary
 Local AS               : 40100          BGP Router Identifier  : 4.4.4.4
 Peers                  :6             Log Neighbor Changes   : No
 Cfg. Hold Time         : 180            Cfg. Keep Alive        : 60
```

b. Ensure BGP AS number and remote-as configured in the neighbor sub-command on peer DUT are correct.
c. Verify neighbor is activated under correct address-family.
d. For BGP session over loopback IP addresses, use the update-source loopback command in the neighbor command.

3. Verify event logs for any discrepancy.

Sample event log:

```
2019-12-03T19:13:34.997328+00:00 L4 hpe-routing[27328]: Event|2901|LOG_
INFO|AMM|-|13.13.13.13: Peer
up. vrf-
name: default
2019-12-03T19:15:58.727524+00:00 L4 hpe-routing[27328]: Event|2902|LOG_
INFO|AMM|-|13.13.13.13: Peer
down. error- code: Cease, error-sub-code: Administrative Shutdown. vrf-name:
default
2019-12-03T19:18:22.519681+00:00 L4 hpe-routing[27328]: Event|2918|LOG_
INFO|AMM|-|4.4.4.4: BGP
identifier sent by Peer 13.13.13.13 matches ours. BGP session may not
established. vrf-name: default
```

FAQs

What is BGP?

BGP (RFC 4271) is an interdomain path-vector routing protocol that runs over TCP on port 179. It provides a reliable loop-free routing between different ASes. BGP routers (called speakers) form peerings (also known as neighbor relationships) with other BGP peers. These peerings can either be internal (iBGP: within the same AS) or external, (eBGP: connecting two different ASs).

What is difference between eBGP and iBGP Peers?

Internal BGP (iBGP): Sessions established with a BGP router that are in the same AS.

External BGP (eBGP): Sessions established with a BGP router that are in a different AS.

How to configure BGP in AOS-CX?

Configuring BGP in AOS-CX involves the following high-level steps:

1. Define AS number
2. Define neighbor specific parameters (Neighbor IP address and Remote AS)
3. Define address families
4. Activate under the sub-address families

How can I configure eBGP neighbors, when eBGP neighbors are directly not connected?

To configure eBGP neighbors when they are not directly connected, you can use the eBGP multihop feature. Multihop configuration works only for eBGP.

```
router bgp <AS>
neighbor <IP-ADDRESS> remote-as <AS-NUMBER>
neighbor <IP-ADDRESS> ebgp-multihop <HOP-COUNT>
```

What is TTL Security in BGP?

TTL Security is a mechanism that protects eBGP peering sessions from attacks caused by forged IP packets. This feature compares the TTL value of an incoming IP packet with a hop-count value configured locally for each eBGP neighbor, and drops the packet if the TTL value is greater than the configured hop count.

```
router bgp <AS>
neighbor ip-address remote-as <Peer AS>
neighbor ip-address ttl-security-hops <number>
```

How can I configure BGP over Tunnel Interfaces?

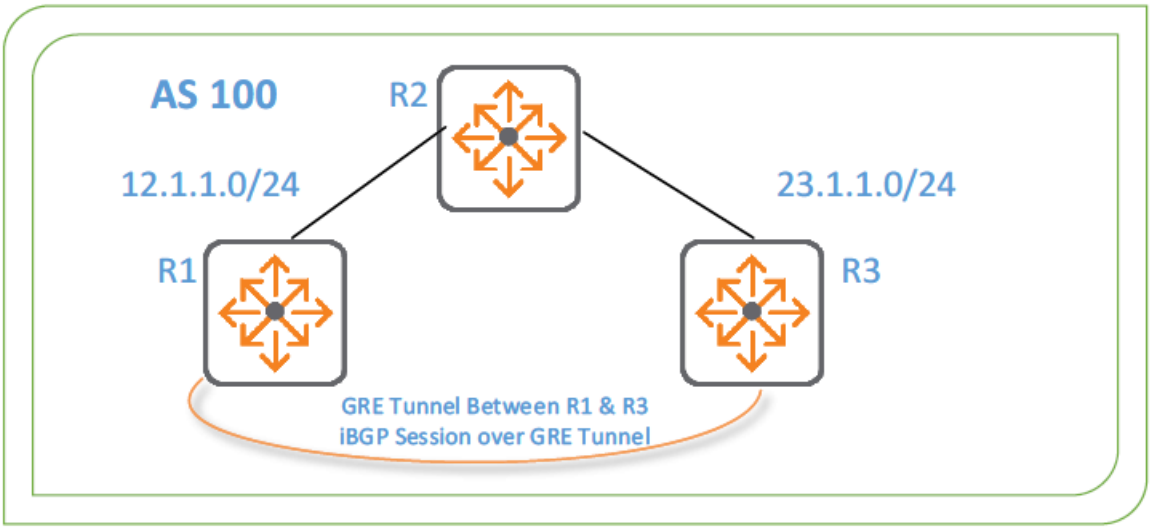

How do I configure iBGP Sessions ?

By default, routes received from an iBGP peer will not be advertised to another iBGP peer. In order to learn routes from another iBGP peer, either full mesh or route-reflector methods can be used.

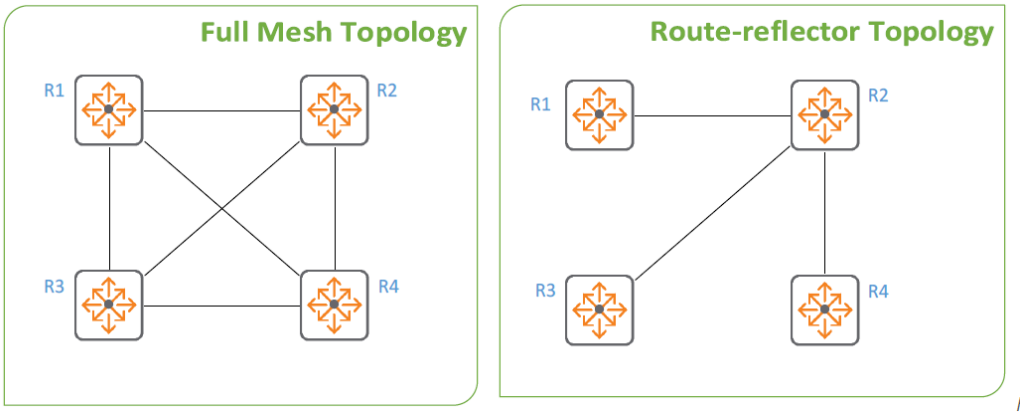

How to advertise prefixes learnt from one iBGP peer to another iBGP Peer Route-Reflectors?

RFC 1966 specifies that IBGP peering can be configured so that it reflects routes to another IBGP peer. The router reflecting route is known as a route reflector (RR) and the router receiving reflected routes is a route reflector client. Three basic rules apply to route reflectors and route reflection.

1. If an RR receives an NLRI from a non-RR client, the RR advertises the NLRI to an RR client. It does not advertise the NLRI to a non-route reflector client.
2. If an RR receives an NLRI from an RR client, it advertises the NLRI to RR and non-RR clients. Even the RR client that sent the advertisement receives a copy of the route, but it discards the NLRI. I
3. If an RR receives a route from an EBGP peer, it advertises the route to RR and non-RR clients.

How do I filter prefixes using Route-Maps ?

You can use route-maps to filter prefixes in BGP by defining specific conditions that determine which prefixes will be permitted or denied in the BGP routing table. The following is a basic example of how to filter prefixes using a route-map:

```
!
ip prefix-list network1 seq 10 permit 192.168.1.0/24 ge 24 le 24
!
route-map 1 permit seq 10
match ip address prefix-list network1
set local-preference 1000
route-map 1 permit seq 20
set local-preference 9000
!
router bgp 10
neighbor 100.1.1.2 remote-as 20
address-family ipv4 unicast
neighbor 100.1.1.2 activate
neighbor 100.1.1.2 route-map 1 in
exit-address-family
```

What are BGP Path Attributes ?

BGP uses path attributes to describe the characteristics of a BGP route. Path attributes are added to BGP route updates and are used by BGP routers to make decisions about the best path to a destination network.

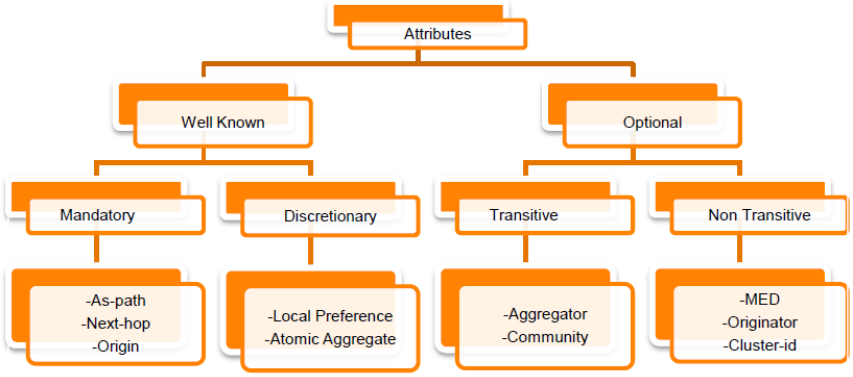

What is BGP Best-Path Selection Algorithm ?

When a BGP router learns multiple routes to the same NLRI, it must choose a single best route. The following list defines the BGP route selection decision process.

1. **Highest administrative weight**: Administrative weight can be assigned to each NLRI locally on a router and the value cannot be communicated to another router. A higher value indicates a better route.
2. **Highest LOCAL_PREF PA**: This attribute can be set on a router inside an AS and distributed inside the AS. A higher value indicates a better route.
3. **Locally injected routed**: Pick the route injected into BGP locally using the network command `redistribution`.
4. **Shortest AS_PATH length**: A shorter AS_PATH indicates a better route. This attribute counts each ASN in the AS_SEQUENCE as one.
5. **Origin PA**: IGP (I) routes are preferred over EGP (E) routes, which are in turn preferred over incomplete (?) routes.
6. **Smallest Multi-Exit Discriminator (MED) PA**: A smaller value indicates a better route.
7. **Neighbor Type**: External BGP (eBGP) routes are better than internal BGP (iBGP) routes.
8. **IGP metric for reaching the NEXT_HOP**: When comparing IGP metrics for each NLRIs NEXT_ HOP, a smaller value indicates a better route.
9. **Keep the oldest eBGP route**: If the routes being compared are eBGP and one is the best path, retain the existing best path. This action reduces eBGP route flaps.
10. **Choose the smallest neighbor router ID**: Use the route with the smallest next-hop router router-id. Only perform this step if bgp bestpath compare-routerid is configured.
11. **Smallest neighbor ID**: The local router has at least two neighbor relationships with a single other router. For this case, the router prefers the route advertised by the lowest neighbor ID.

How do I advertise Prefixes using Network Command?

The network command advertises the prefix only if the same prefix and mask is present in the routing table whether it is connected, static or learned dynamically via any routing protocol.

```
network 192.168.1.0/24 [route-map <NAME>]
```

How do I redistribute prefixes into BGP?

To redistribute prefixes into BGP, follow these steps:

1. Configure the BGP instance with the AS number.
2. Specify routes to import into BGP. This command causes routes from the specified protocol to be considered for redistribution into BGP.
3. If required, specify a route map to apply to the redistributed route.

Rredistribute has following options under router bgp sub-address family:

`connected`: Redistribute directly attached networks

`ospf`: Redistribute OSPF routes (both OSPFv2 and OSPfv3)

`static`: Redistribute static routes

How do I advertise a default-route into BGP?

Use the default-originate command to advertise a default-route into the neighbor. When default-originate is configured on a neighbor, then this router will be used as default exit point for that neighbor.

```
neighbor <IP-ADDRESS> default-originate [route-map <NAME>]
```

How do I aggregate routes in BGP?

Use the aggregate-address command to aggregate more specific routes into one route. You can minimize the size of the routing table using this command.

```
aggregate-address <IP-ADDR>/<MASK> [as-set] [summary-only]
       [suppress-map <MAP-NAME>] [advertise-map <MAP-NAME>]
       [attribute-map <MAP-NAME>]
```

How do I advertise the aggregated route back to advertising peers?

On 8360 Switch series, you can use the [no] aggregate-address split-horizon command to enable or disable the split horizon rule for advertising the aggregated summary route. When BGP learns a network from a per, and this is part of the locally aggregated route, then BGP will not advettise the aggregated route to this BGP peer. The aggregate-address split-horizon feature is enabled by default and will not show up in the running configuration.

What are the benefits of configuring BGP peer groups?

Configuring BGP peer groups provides several benefits, particularly for iBGP peers where a common set of sub-commands (such as remote-as, route-map, RR-client, update-source, and route-map) are repeated for all peers.

Does the iBGP split-horizon rule does apply across address families?

On 8325, 8325P and 8325H Switch series, iBGP split-horizon rule does not apply across address families For example, routes exchanged between VSX peers over iBGP may be re-advertised into other families, which can introduce routing loops. To prevent this behavior, route-maps (for example, using BGP communities) should be used to classify the routes and explicitly control their propagation. In particular, it is recommended to enforce policies using route-maps that restrict the advertisement of these routes into the EVPN address family.

How do I configure MP-BGP?

To configure MP-BGP (Multiprotocol BGP), follow these steps:

1. Configure the router with a unique router ID.
2. Enable the address family for the network you want to advertise.
3. Configure the neighbor with the remote AS number.
4. Configure other address families for the neighbor to enable MP-BGP support.

Sample Configuration:

```
router bgp 6488163
  bgp router-id 1.1.1.1
  neighbor 1000::2 remote-as 6422626
  address-family ipv4 unicast
    neighbor 1000::2 activate
     neighbor 1000::2 send-community extended
   exit-address-family
   address-family ipv6 unicast
   neighbor 1000::2 activate
    neighbor 1000::2 send-community extended
```

How to use regular expression in BGP prefix filtering?

Regular expressions can be used with AS-path lists or extended community lists to match AS numbers or communities, respectively. AS-path lists are used to filter inbound or outbound BGP route updates based on the AS path attribute. To use regular expressions in AS-path lists, you can use a pattern matching method similar to Unix style pattern matching to match the AS numbers in the AS path attribute.

What is next-hop address tracking?

BGP next-hop address tracking is a feature that improves convergence time by quickly notifying BGP of next-hop changes as they are updated in the RIB. Though this functionality improves the BGP convergence time, this needs to be cautiously enabled as this might cause the instability in the network if the link which connects to the BGP peer flaps so frequently.

What is BGP route-flap dampening?

BGP route-flap dampening is a mechanism used to prevent BGP from selecting unstable routes. A route flap occurs when a BGP route changes state frequently, from up to down or vice versa.

What are BGP standard error codes?

Error Code Value	Code Name
1	Message Header Error
2	Open Message Error
3	Update Message Error
4	Hold Timer Expired
5	Finite State Machine Error
6	Cease

What are the common misconfigurations that should be avoided?

The following are some common misconfigurations:

- Missing neighbor activate command on remote DUT
- Wrong remote-as Number configured for neighbor
- Same Router-ID enabled on both peers
- Path MTU mismatch
- `update-source` command missing for multi-hop session
- send community missing at neighbor levels

When do I need to use the `update source loopback` command?

You will need to use the `update-source loopback` command if the IP address you are using for the `neighbor` command in BGP is not directly attached to an interface on your device. In such cases, you must use the update-source command along with loopback, ip-address, or VLAN to specify the source IP address for the TCP connection.

```
neighbor ip-address update-source loopback
```

BGP commands

address-family

```
address-family {{ipv4 | ipv6} unicast | l2vpn evpn}
no address-family {{ipv4 | ipv6} unicast | l2vpn evpn}
```

Description

Specifies address family to use and changes to the configuration context for the specified family:

- **config-bgp-ipv4-uc** for IPv4 unicast
- **config-bgp-ipv6-uc** for IPv6 unicast
- **config-bgp-l2vpn-evpn** for L2VPN EVPN

The **no** form of this command removes the specified address family configuration.

Parameter	Description
`ipv4`	Selects the IPv4 address family.
`ipv6`	Selects the IPv6 address family.
`unicast`	Specifies unicast addresses.
`l2vpn evpn`	Selects the L2VPN EVPN address family. Route maps with the **match vni** clause can be used with L2VPN EVPN neighbors only.

Example

Setting the address family to IPv4 unicast.

```
switch(config-bgp)# address-family ipv4 unicast
switch(config-bgp-ipv4-uc)#
```

Setting the address family to L2VPN EVPN.

```
switch(config-bgp)# address-family l2vpn evpn
switch(config-bgp-l2vpn-evpn)#
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

aggregate-address

```
aggregate-address <IP-ADDR>/<MASK> [as-set] [summary-only]
        [suppress-map <MAP-NAME>] [advertise-map <MAP-NAME>]
        [attribute-map <MAP-NAME>]

no aggregate-address <IP-ADDR>/<MASK> [as-set] [summary-only]
        [suppress-map <MAP-NAME>] [advertise-map <MAP-NAME>]
        [attribute-map <MAP-NAME>]

[no] aggregate-address split-horizon
```

Description

Creates an aggregate address entry in the BGP routing table.

The **no** form of this command removes the specified aggregate address entry.

Parameter	Description
<ADDRESS>	Specifies an aggregate address in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255, or IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a

Parameter	Description
	hexadecimal number from 0 to F.
<MASK>	Specifies the number of bits in the address mask in CIDR format (**x**), where **x** is a decimal number from 0 to 128.
as-set	The AS_PATH attribute advertised for this route will contain an AS_SET consisting of all AS numbers contained in all paths that are being summarized.
summary-only	Creates the aggregate route but also suppresses advertisements of more-specific routes to all neighbors.
suppress-map *<MAP-NAME>*	Specifies an aggregate route for creation, but suppresses the advertisement of the created route. Match clauses of route maps can be used to suppress some more-specific routes of the aggregate selectively, and leave others unsuppressed. IP prefix lists and as_path lists match clauses are supported.
advertise-map *<MAP-NAME>*	Specifies routes that will be used to build attributes of the aggregate route, such as AS_SET or community.
attribute-map *<MAP-NAME>*	Specifies that the attributes of the aggregate route can be changed.
split-horizon	(On 8360 Switch series only) Disables advertisement of an aggregated peer.

Examples

```
switch(config-bgp-ipv4-uc)# aggregate-address 10.0.0.0/8
switch(config-bgp-ipv4-uc)# no aggregate-address 10.0.0.0/8
```

```
switch(config-bgp-ipv6-uc)# aggregate-address 2001:0db8:85a3::8a2e:0370:7334/24
switch(config-bgp-ipv6-uc)# no aggregate-address 2001:0db8:85a3::8a2e:0370:7334/24
```

```
switch(config-bgp-ipv4-uc)# aggregate-address 10.0.0.0/8 as-set summary-only
switch(config-bgp-ipv4-uc)# aggregate-address 10.0.0.0/8 attribute-map RMap
```

```
switch(config-bgp-ipv6-uc)# aggregate-address 2001:0db8:85a3::8a2e:0370:7334/24
as-set summary-only
switch(config-bgp-ipv6-uc)# aggregate-address 2001:0db8:85a3::8a2e:0370:7334/24
attribute-map RMap
```

Command History

Release	Modification
10.14	The **split-horizon** parameter is introduced on 8360 switch series.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp-ipv4-uc` `config-bgp-ipv6-uc`	Administrators or local user group members with execution rights for this command.

bgp add-paths select best 2

```
bgp add-paths select best 2
no bgp add-paths select best 2
```

Description

Enables BGP-PIC and is available in the router BGP context.

The **no** form of this command disables BGP-PIC.

Parameter	Description
`additional-paths`	Specifies the number of paths.

Usage

- When BGP-PIC is enabled or disabled, there will be a momentary traffic loss for a few seconds due to all BGP routes being potentially deleted and re-added as necessary.
- Two paths are required for protected routes.

Examples

Enabling BGP-PIC:

```
switch(config)# router bgp 1
switch(config-bgp)# address-family ipv4 unicast
switch(config-bgp-ipv4-uc)# bgp add-paths select best 2
switch(config-bgp-ipv4-uc)#
```
```

Disabling BGP-PIC:

```
switch(config)# router bgp 1
switch(config-bgp)# address-family ipv4 unicast
switch(config-bgp-ipv4-uc)# no bgp add-paths select best 2
switch(config-bgp-ipv4-uc)#
```
```

Command History

Release	Modification
10.14.1000 or later	Command introduced on 9300S Switch Series.

Command Information

Platforms	Command context	Authority
9300 9300S	`config-bgp`	Administrators or local user group members with execution rights for this command.

bgp always-compare-med

```
bgp always-compare-med
no bgp always-compare-med
```

Description

Enables comparison of the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems. Any changes in BGP configuration are applied by restarting the current BGP sessions on the VRFs.

The **no** form of this command sets comparison of MED to the default setting (disabled).

Usage

- MED is one of the parameters that is considered when selecting the best path among many alternative paths. The path with a lower MED is preferred over a path with a higher MED.
- During the best-path selection process, MED comparison is done only among paths from the same autonomous system. Use the command `bgp always-compare-med` to change this behavior by enforcing MED comparison between all paths, regardless of the autonomous system from which the paths are received.

Examples

```
switch(config-bgp)# bgp always-compare-med
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

```
switch(config-bgp)# no bgp always-compare-med
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

bgp asnotation dotted

```
bgp asnotation dotted
no bgp asnotation dotted
```

Description

Specifies that Autonomous System (AS) numbers greater than 65535 be shown in dotted integer format for all show commands, including running-configuration.

The **no** form of this command restores the default format of non-dotted, simple integer.

Example

```
switch(config-bgp)# bgp asnotation dotted
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

bgp asnotation dotted-plus

```
bgp asnotation dotted-plus
no bgp asnotation dotted-plus
```

Description

Specifies that all Autonomous System (AS) numbers be shown in dotted integer format for all show commands, including running-configuration.

The **no** form of this command restores the default format of non-dotted, simple integer.

Example

```
switch(config-bgp)# bgp asnotation dotted-plus
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-bgp	Administrators or local user group members with execution rights for this command.

bgp bestpath as-path ignore

```
bgp bestpath as-path ignore
no bgp bestpath as-path ignore
```

Description

Configures BGP to avoid considering the autonomous system (AS) path during best path route selection. By default, the AS-path is considered during BGP best path selection.

Any changes in BGP configuration are applied by restarting the current BGP sessions on the VRFs.

The **no** form of this command restores default behavior which configures BGP to consider the AS-path during route selection.

Examples

```
switch(config-bgp)# bgp bestpath as-path ignore
All current BGP sessions in VRF default will be restarted.
```

```
Do you want to continue (y/n)?
```

```
switch(config-bgp)# no bgp bestpath as-path ignore
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-bgp	Administrators or local user group members with execution rights for this command.

bgp bestpath as-path multipath-relax

```
bgp bestpath as-path multipath-relax
no bgp bestpath as-path multipath-relax
```

Description

Configures Border Gateway Protocol (BGP) to treat two BGP routes as equal cost even if their AS-paths differ, as long as their AS-path lengths and other relevant attributes are the same. This allows routes with different AS-paths to be programmed into the forwarding table as equal cost multipath routes.

Any changes in BGP configuration are applied by restarting the current BGP sessions on the VRFs.

The **no** form of this command restores the default behavior which configures BGP to treat two BGP routes as different costs when their AS-paths differ.

Examples

```
switch(config-bgp)# bgp bestpath as-path multipath-relax
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

```
switch(config-bgp)# no bgp bestpath as-path multipath-relax
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-bgp	Administrators or local user group members with execution rights for this command.

bgp bestpath bgp-std-selection

```
bgp bestpath bgp-std-selection
no bgp bestpath bgp-std-selection
```

Description

Configures the standard Border Gateway Protocol (BGP) best-path selection process used by major vendors. This command is applied in router configuration mode and is VRF-specific. When enabled, it modifies how redistributed routes are considered during BGP best-path selection.

The **no** form of this command restores the default behavior—HPE-ANW-specific best-path selection. This resets all active BGP peers and may change routing path selection across the network for all BGP peers within the VRF where the command is configured.

Examples

Configuring BGP best-path selection:

```
switch(config-bgp)# bgp bestpath bgp-std-selection
All active BGP sessions in the VRF %s will be restarted and also may change
the routing path selection in the network for all BGP peers in the VRF.
Do you want to continue (y/n)?
```

Restoring to default behavior:

```
switch(config-bgp)# no bgp bestpath bgp-std-selection
All active BGP sessions in the VRF %s will be restarted and also may change
the routing path selection in the network for all BGP peers in the VRF.
Do you want to continue (y/n)?
```

Command History

Release	Modification
10.17	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	

bgp bestpath compare-routerid

```
bgp bestpath compare-routerid
no bgp bestpath compare-routerid
```

Description

Configures a BGP routing process to compare identical routes received from different external peers during the best path selection process and selects the route with the lowest router ID as the best path. Defaults to disabled.

Any changes in BGP configuration are applied by restarting the current BGP sessions in the VRFs.

The **no** form of this command returns the BGP routing process to the default operation. By default, BGP selects the route that was received first when two routes with identical attributes are received.

Examples

```
switch(config-bgp)# bgp bestpath compare-routerid
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

```
switch(config-bgp)# no bgp bestpath compare-routerid
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

bgp bestpath med confed

```
bgp bestpath med confed
no bgp bestpath med confed
```

Description

Compares the identical routes received from the different confederation peers and selects the route with the lowest Multi Exit Discriminator (MED) value as the best path. This behavior is disabled by default.

The **no** form of this command prevents the routing process from considering the MED value.

The selection of other attributes like `as-multi-path relax` and `as-path ignore` will not affect the behavior of this command within a confederation.

Examples

Selecting the route with lowest MED value:

```
switch(config-bgp)#  bgp bestpath med confed
All active BGP sessions in the VRF %s will be restarted.
Do you want to continue (y/n)?
```

Preventing the routing process from selecting the MED value:

```
switch(config-bgp)# no bgp bestpath med confed
All active BGP sessions in the VRF %s will be restarted.
Do you want to continue (y/n)?
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

bgp bestpath med missing-as-worst

```
bgp bestpath med missing-as-worst
no bgp bestpath med missing-as-worst
```

Description

Configures a BGP routing process to assign a value of infinity (max possible) to routes that are missing the Multi Exit Discriminator (MED) attribute. The path without a MED value is the least desirable path. Any changes in BGP configuration are applied by restarting the current BGP sessions in the VRFs.

The **no** form of this command restores default behavior. The default behavior assigns a value of 0 to the missing MED.

Examples

```
switch(config-bgp)# bgp bestpath med missing-as-worst
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

```
switch(config-bgp)# no bgp bestpath med missing-as-worst
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

bgp cluster id

```
bgp cluster-id {<IPV4-ADDR> | <ID>}
no bgp cluster-id {<IPV4-ADDR> | <ID>}
```

Description

Specifies the cluster ID when the BGP router is used as a route-reflector. The cluster ID default is the router ID. Any changes in BGP configuration are applied by restarting the current BGP sessions on the VRFs.

The **no** form of this command sets the cluster ID to the default value, which is the router ID.

Parameter	Description
<IPV4-ADDR>	Specifies an IP address in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. You can remove leading zeros. For example, the address **192.169.005.100** becomes **192.168.5.100**.
<ID>	Specifies the cluster ID as 32-bit number. Range: 1 to 4294967295.

Examples

```
switch(config-bgp)# bgp cluster-id 2.2.2.2
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

```
switch(config-bgp)# no bgp cluster-id 2.2.2.2
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

bgp confederation

```
 bgp confederation <AS-NUMBER>
no bgp confederation<AS-NUMBER>
```

Description

Configures a BGP confederation with the confederation identifier. The group of Autonomous Systems (ASs) will be presented as a single autonomous system with the confederation identifier as the AS number.

The **no** form of the command deletes the BGP confederation identifier.

Parameter	Description
<AS-NUMBER>	Sets the identifier for the confederation. Range:1-4294967295.

Examples

Configuring the BGP confederation with the AS number:

```
switch(config-bgp)# bgp confederation 100
```

Deleting BGP confederation identifier:

```
switch(config-bgp)# no bgp confederation 100
This will delete BGP confederation identifier on this device.
Do you want to continue (y/n)?
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

bgp confederation peers

```
bgp confederation peers <AS NUMBER>
no bgp bgp confederation peers <AS NUMBER>
```

Description

Configures BGP confederation peers with both same and different sub-autonomous system to establish an eBGP membership. You can configure a list of AS numbers separated by spaces.

The **no** form of this command disables the peer session and deletes the peer information.

Parameter	Description
<AS NUMBER>	Specifies the autonomous system numbers to establish an eBGP membership. Range: 64512-65535.

Examples

```
switch(config-bgp)# bgp confederation peers 64512 64513
```

```
switch(config-bgp)# no bgp confederation peers 64512
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A,	`config-bgp`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000		

bgp dampening

```
bgp dampening {{half-life <HALF-TIME> reuse <LOW-THRESHOLD> suppress <HI-THRESHOLD> max-
suppress-time <MAX-TIME>}
| route-map <NAME>}
```

Description

Enables route flap dampening which reduces the propagation of unstable routes in the network.

Parameter	Description
`half-life <HALF-TIME>`	Specifies the half-life time in minutes. When the time expires, the penalty on a route gets reduced exponentially to half its current value. Default: 15.
`reuse <LOW-THRESHOLD>`	Specifies the lower threshold of penalty. On a suppressed route, when the penalty on a route falls below this value, the route is unsuppressed. Default: 750.
`suppress <HI-THRESHOLD>`	Specifies the upper threshold of penalty. When the penalty on a flapping route exceeds this value, the route is suppressed. Default: 2000.
`max-suppress-time <MAX-TIME>`	Specifies the maximum time to keep a route suppressed in minutes. Once this timer expires, the route is unsuppressed. Default: 60.
`route-map <NAME>`	Specifies the name of a route map.

These parameters can be configured at the router level for specific address families or the same parameters can be configured under a route map which can be applied to dampening command.

Usage

The dampening algorithm assigns a penalty of 1000 to a flapping route every time the route gets withdrawn. The penalty values accumulate on the route every time it flaps. However, the penalty decays and is reduced to half its value by the half-life time.

This feature is not applicable on IBGP routes.

Example

```
switch(config)# router bgp 1
switch(config-bgp)# address-family ipv4 unicast
switch(config-bgp-ipv4-uc)# bgp dampening

switch(config-bgp-ipv4-uc)# bgp dampening route-map abc

switch(config-bgp-ipv4-uc)# bgp dampening route-map xyz

switch(config-bgp-ipv4-uc)# bgp dampening half-life 10 reuse 100 suppress 250 max-
suppress-time 45
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp-ipv4-uc`	Administrators or local user group members with execution rights for this command.

bgp default local-preference

```
bgp default local-preference <NUMBER>
no bgp default local-preference
```

Description

Default local preference value for BGP learned routes. Any changes in BGP configuration are applied by restarting the current BGP sessions on the VRFs.

The **no** form of this command sets the local preference to the default value of 100.

Parameter	Description
<NUMBER>	Specifies the local preference value. Range: 0 to 4294967295. Default: 100.

Examples

```
switch(config-bgp)# bgp default local-preference 20
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

```
switch(config-bgp)# no bgp default local-preference
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

bgp deterministic-med

```
bgp deterministic-med
no bgp deterministic-med
```

Description

Enables comparison of the Multi-Exit Discriminator (MED) attribute when selecting routes advertised by different peers in the same autonomous system. Any changes in BGP configuration are applied by restarting the current BGP sessions on the VRFs.

The **no** form of this command sets MED comparison to the default setting of disabled.

Examples

```
switch(config-bgp)# bgp deterministic-med
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

```
switch(config-bgp)# no bgp deterministic-med
All current BGP sessions in VRF default will be restarted.
```

```
Do you want to continue (y/n)?
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

bgp fast-external-fallover

```
bgp fast-external-fallover
no bgp fast-external-fallover
```

Description

Sets the switch to reset the BGP sessions of any directly adjacent external peers when the connected link goes down. It is enabled by default.

The **no** form of this command restores the default behavior where BGP waits until the hold time expires before closing sessions.

Examples

```
switch(config-bgp)# bgp fast-external-fallover
switch(config-bgp)# no bgp fast-external-fallover
```

Command History

Release	Modification
10.08	The default behavior has been changed from disabled to enabled state. **NOTE:** When upgrading, the feature will remain in the state it was

Release	Modification
	(disabled or enabled) in the earlier release.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

bgp graceful-restart restart-time

```
bgp graceful-restart restart-time <DELAY>
no bgp graceful-restart restart-time
```

Description

Sets the graceful restart timer which determines how long the switch waits for a graceful-restart capable neighbor to re-establish BGP peering. Any changes in BGP configuration are applied by restarting the current BGP sessions on the VRFs.

The **no** form of this command resets to the default value of 120 seconds.

Parameter	Description
<DELAY>	Graceful restart timer delay in seconds. Range: 1 to 3600. Default: 1500.

Usage

- Graceful restart functionality is enabled by default, and there is no command to disable the functionality at the protocol level.
- However, the graceful-restart functionality can be disabled globally using the command **router graceful-restart**.

Examples

```
switch(config-bgp)# bgp graceful-restart restart-time 150
All current BGP sessions in VRF default will be restarted.
```

```
Do you want to continue (y/n)?
```

```
switch(config-bgp)# no bgp graceful-restart restart-time
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-bgp	Administrators or local user group members with execution rights for this command.

bgp graceful-restart stalepath-time

```
bgp graceful-restart stalepath-time <TIME>
no bgp graceful-restart stalepath-time
```

Description

Sets the stale path timer. This timer determines how long BGP keeps stale routes from the restarting BGP peer. Any changes in BGP configuration are applied by restarting the current BGP sessions on the VRFs.

The **no** form of this command resets to the stale path timer to the default of 300 seconds.

Parameter	Description
<TIME>	Specifies the stale path timer in seconds. Range: 1 to 3600. Default: 300.

Examples

```
switch(config-bgp)# bgp graceful-restart stalepath-time 300
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

```
switch(config-bgp)# no bgp graceful-restart stalepath-time
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-bgp	Administrators or local user group members with execution rights for this command.

bgp log-neighbor-changes

```
bgp log-neighbor-changes
no bgp log-neighbor-changes
```

Description

Enables logging of BGP neighbor session state changes.

The **no** form of this command disables logging of changes in BGP neighbor adjacencies.

Examples

```
switch(config-bgp)# bgp log-neighbor-changes
switch(config-bgp)# no bgp log-neighbor-changes
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

bgp maxas-limit

```
bgp maxas-limit <LENGTH>
no bgp maxas-limit
```

Description

Specifies the maximum size of AS paths in update messages. Routes with AS paths greater than the specified length are discarded.

The **no** form of this command sets the limit to the default of 32.

Parameter	Description
<LENGTH>	Specifies the number of AS segments. Length: 1 to 32 characters. Default: 32.

Example

```
switch(config-bgp)# bgp maxas-limit 20
switch(config-bgp)# no bgp maxas-limit
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325	`config-bgp`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325H 8325P 8360 9300 9300S 10000		

bgp router-id

```
bgp router-id <ROUTER-ID>
no bgp router-id <ROUTER-ID>
```

Description

Configures a fixed router ID for the BGP peer process running on the router. Any changes in BGP configuration are applied by restarting the current BGP sessions on the VRFs.

The **no** form of this command removes the fixed router ID from the running configuration and restores the default router ID selection.

Parameter	Description
<ROUTER-ID>	Specifies the router ID in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255. If router-id is changed, then all the active BGP peer sessions go down and restart with the newly configured router-id.

Usage

BGP determines the router ID as follows:

1. The address configured with the command `bgp router-id`.
2. The highest IP address on all the loopback interfaces.
3. The highest IP address on any interface.

Examples

```
switch(config-bgp)# bgp router-id 1.1.1.1
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

```
switch(config-bgp)# no bgp router-id 1.1.1.1
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

bgp update-group

```
bgp update-group
no bgp update-group
```

Description

Enables the update-group feature. Update-group is used to optimize BGP processing. With update-group enabled, neighbors that share same outbound policy are grouped under single update-group and update messages are shared to all peers within that group. Disabled by default.

On enabling/disabling update-group, all the current BGP sessions including dynamic peer sessions will be restarted. The dynamic peer sessions will get re-established based on the connect-retry timer configured on the peer device.

BGP determines the AS number size to be used with each individual peer based on the configured AS octet-size on the peers. If BGP is configured to use 2 octet ASN with peer A and 4 octet ASN with peer B, then BGP updates sends to peer A and peer B will be different. Therefore, they will form into two different update-groups.

This feature is supported for IBGP alone. It is not recommended to enable update-group on VPNv4 address families or configure update-group if an VPNv4 address family is configured. Update-group indices are created automatically and cannot be configured. Indices are ephemeral in nature. Update-group indices will not be retained across process restart. When update-group is configured, route reflector clients and non route reflector clients are parted into different update-groups. The use of update-groups imposes the following limitations in BGP processing:

1. While advertising routes to peers of the update-group, the check on route's next-hop is ignored. This means that, BGP routes can get advertised to a peer even if the next-hop IP of the routes is same as peer's IP, and the peer should discard such routes.
2. The use of ORF is restricted. BGP will continue to send outbound route filters to peers, but received filters are ignored.
3. Sender loop detection is delayed until the point where UPDATE messages are sent to the members of the update-group. This means route dropped by the sender due to loop detection will still count as advertised, as they have been advertised to the update-groups.

It is recommended to enable BGP update-group globally before configuring BGP neighbors. This ensures that toggling of BGP neighborship can be avoided and memory of the routing process can be optimized.

The **no** form of this command disables the update-group feature.

Examples

Enabling BGP update-group:

```
switch(config-bgp)# bgp update-group
All current BGP sessions in VRF default will be restarted.
 Do you want to continue (y/n)?
```

Disabling BGP update-group:

```
switch(config-bgp)# no bgp update-group
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)?
```

Command History

Release	Modification
10.13	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

clear bgp

```
clear bgp [vrf <VRF-NAME>][ip unicast | ipv6 unicast | l2vpn evpn | all]
     {* | <NEIGHBOR-IP-ADDRESS> | <PEER-GROUP-NAME> } [soft {in | out} ]
```

Description

Clears BGP with a hard or soft reset. If the address-family is specified, then the neighbors that negotiated the respective family will be reset or a route refresh message will be sent.

A route refresh message is sent when **soft in** or **soft out** is specified.

Parameter	Description
`ip unicast`	Specifies the IPv4 address family.

Parameter	Description
`ipv6 unicast`	Specifies the IPv6 address family.
`l2vpn evpn`	Selects the L2VPN EVPN address family
`vrf <VRF-NAME>`	Specifies a VRF name. By default, the default BGP instance information is displayed if the VRF is not specified.
`all`	Specifies all VRFs and address families.
`* | <NEIGHBOR-IP-ADDRESS> | <PEER-GROUP-NAME>`	Specifies a neighbor IP address or peer group name for which peer sessions are to be reset, or `*` to reset all sessions.
`soft in`	Clears the BGP in the inbound direction without a hard reset of the session. This will send a route refresh request message.
`soft out`	Clears the BGP in the outbound direction without a hard reset of the session. This will send a route refresh message.

Examples

```
switch(config)# clear bgp all *
```

```
switch# clear bgp all PEER_GRP_1 soft in
switch# clear bgp all PEER_GRP_1 soft out
```

```
switch# clear bgp ip unicast 9.0.0.2 soft in
switch# clear bgp ip unicast 9.0.0.2 soft out
switch# clear bgp ip unicast PEER_GRP_1 soft in
switch# clear bgp ip unicast PEER_GRP_1 soft out
```

```
switch# clear bgp vrf v1 ip unicast 9.0.0.2 soft in
switch# clear bgp vrf v1 ip unicast 9.0.0.2 soft out
switch# clear bgp vrf v1 ip unicast PEER_GRP_1 soft in
switch# clear bgp vrf v1 ip unicast PEER_GRP_1 soft out
switch# clear bgp vrf v1 all PEER_GRP_1 soft in
switch# clear bgp vrf v1 all PEER_GRP_1 soft out
```

```
switch# clear bgp ipv6 unicast 9000::2 soft in
switch# clear bgp ipv6 unicast 9000::2 soft out
switch# clear bgp vrf v1 ipv6 unicast 9000::2 soft in
switch# clear bgp vrf v1 ipv6 unicast 9000::2 soft out
```

```
switch# clear bgp l2vpn evpn * soft in
switch# clear bgp l2vpn evpn * soft out
switch# clear bgp l2vpn evpn 9.0.0.2 soft in
```

```
switch# clear bgp l2vpn evpn 9.0.0.2 soft out
switch# clear bgp l2vpn evpn PEER_GRP_1 soft in
switch# clear bgp l2vpn evpn PEER_GRP_1 soft out
```

Command History

Release	Modification
10.15	**soft out** and **<*PEER-GROUP-NAME*> soft in** parameters are added.
10.14	Replaced IPv4 with IP in commands. IPv4 is deprecated.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

disable|enable

```
disable | enable
```

Description

This command disables or enables the BGP instance while retaining the configuration. Disable and enable of the BGP instance may result in a change of the router ID.

By default the BGP instance is enabled.

Examples

```
switch(config)# router bgp 100
switch(config-bgp)# disable
```

The routes learned from a peer will not be removed immediately from the BGP's local RIB when the BGP session on the peer device is disabled. The peer routes experience a graceful restart.

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

distance bgp

```
distance bgp <EXTERNAL> <INTERNAL>
no distance bgp <EXTERNAL> <INTERNAL>
```

Description

Configures the administrative distance for BGP.

The **no** form of this command restores the default settings, 20 for eBGP and 200 for iBGP,

Parameter	Description
<EXTERNAL>	Specifies the administrative distance for eBGP routes. Range: 1 to 255. Default: 20.
<INTERNAL>	Specifies the administrative distance for iBGP routes. Range: 1 to 255. Default: 200.

Example

```
switch(config-bgp-ipv4-uc)# distance bgp 100 150
switch(config-bgp-ipv4-uc)# no distance bgp 100 150
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp-ipv4-uc` `config-bgp-ipv6-uc`	Administrators or local user group members with execution rights for this command.

maximum-paths

```
maximum-paths <MAXPATHS>
no maximum-paths <MAXPATHS>
```

Description

Configures the maximum number of paths that BGP adds to the route table for equal-cost multipath (ECMP) load balancing for routes learned from both internal and external BGP. Any changes in BGP configuration are applied by restarting the current BGP sessions on the VRFs.

On HPE Aruba Networking 8325, 8325H, 8325P and 10000 switch series, a maximum-paths configuration is supported globally as well as for **ipv4-unicast**, **ipv6-unicast** and **l2vpn-evpn** address families. For HPE Aruba Networking 8325, 8325H, and 10000 switch series, this configuration is supported under the **l2vpn-evpn** for L2/L3 Overlay ECMP feature. On HPE Aruba Networking 6300, 6400, 8100 and 8360 switch series, this command is supported under the **l2vpn-evpn** for L2 Overlay ECMP feature.

If address-family specific maximum-paths are configured, they take precedence over global configuration. Greater than eight maximum-paths at can be configured within the global context, but for **l2vpn-evpn** routes the supported maximum limit is still eight.

The **no** form of this command restores the default setting of 4.

For VxLAN L2 Overlay ECMP: On HPE Aruba Networking 6300, 6400, 8100, 8325, 8360 and 10000, this command also controls the maximum number of paths that BGP adds to the mac-table for equal-cost multipath (ECMP) load balancing for VxLAN overlay mac routes, reachable via VXLAN tunnels.

Parameter	Description
<MAXPATHS>	Specifies the maximum number of paths. 1 to 32 paths can be set globally. 1 to 8 paths can be set for **l2vpn-evpn** address family.1 to 32 paths can be set for **ipv4-unicast** and **ipv6-unicast**.

Usage

When both global and address-family maximum paths are configured then the address family value takes precedence.

Global max-path	l2evpn, ipv4-unicast, or ipv6-unicast address-family Max-path	Operational Max-path	Comments
Default(4)	Default (4)	4	--
Configured (4)	configured (4)	4	Four is the configured value for address family paths
Default (4)	8 (max)	8	The address family value takes precedence
32 (max)	8 (max)	8	CLI configuration is limited to eight paths.
Default (8)	configured (4)	4	Four is the configured value for address family paths

EVPN Multi-Homing Interoperability

For HPE Aruba Networking 6300 (except 6300L), 6400, 8100, 8325, 8325H, 8360 and 10000 switch series:

EVPN MAC/IP (Type-2) routes with either non-reserved or reserved ESI values of 0 or MAX-ESI can be evaluated for forwarding. EVPN MAC/IP routes (Type-2 routes) with a reserved ESI value are resolved by the MAC/IP route only, using the BGP next-hop within the type-2 route. EVPN MAC/IP routes with a non-reserved ESI value are resolved only if the route has an associated per-ES Ethernet Auto-Discovery route (a type-1 route).

6300 (except 6300L), 6400, 8100, 8325, 8325H, 8325P, 8360 and 10000 switch series using EVPN MAC/IP route resolution with non-reserved ESI values can use VSX multi-homing for locally connected devices while coexisting in a EVPN-VXLAN fabric with other switches using EVPN multi-homing to connect local devices. The MAC and IP addresses of remote endpoints are learned from those remote switches using EVPN control plane messages, and are assigned multiple next-hop IP addresses (unique VTEP addresses that identify switches using EVPN multi-homing).

Examples

Configuring the maximum number of global paths:

```
switch(config)# router bgp 1
switch(config-bgp)# maximum-paths 32
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)? y
```

Configuring the maximum paths for the IPv4-unicast address family:

```
switch(config-bgp)# address-family ipv4 unicast
switch(config-bgp-ipv4-uc)#  maximum-paths 6
```

Configuring the maximum-paths paths for the IPv6-unicast address family:

```
switch(config-bgp)# address-family ipv6 unicast
switch(config-bgp-ipv6-uc)#  maximum-paths 6
```

Configuring the maximum-paths for the l2vpn-evpn address family:

```
switch(config-bgp)# address-family l2vpn evpn
switch(config-bgp-l2vpn-evpn)# maximum-paths 6
```

Removing the global maximum paths setting:

```
switch(config-bgp)# no maximum-paths
All current BGP sessions in VRF default will be restarted.
Do you want to continue (y/n)? y
```

Command History

Release	Modification
10.16.1000	Adds support to HPE Aruba Networking 8325, 8325H, 10000 switch series for L2 and L3 Overlay ECMP features. On HPE Aruba Networking 6300, 6400, 8100 and 8360 switch series, this command is supported under the **l2vpn-evpn** for L2 Overlay ECMP feature.
10.13.1000	On 8325, 8325H, 8325P and 10000 switch series, configuration is supported on **ipv4-unicast**, **ipv6-unicast** and **l2vpn-evpn** address families. For HPE Aruba Networking 6300, 6400, 8100, 8360, 8400 and 9300 Switch series, this configuration is supported at the global level only.
10.10	Increased upper limit of range of <*MAXPATHS*> parameter to 32.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority

6300, except for S3L75A, S3L76A, S3L77A
6400
8100
8325
8325H
8325P
8360
9300
9300S
10000
	`config-bgp`
	`config-bgp-ipv4-uc`
	`config-bgp-ipv6-uc`
	`config-bgp-l2vpn-evpn`
		Administrators or local user group members with execution rights for this command.

neighbor activate

`neighbor {<IP-ADDR>|<IFNAME>} activate`
`no neighbor {<IP-ADDR>|<IFNAME>}activate`

Description

This command enables the address-family capability and exchange of information specific to an address family with a BGP neighbor.

The **no** form of this command removes the address-family capability and disables the exchange of routes for the specified address-family with the BGP neighbor.

Parameter	Description
<IP-ADDR>	Specifies an IP address in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255, or IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F. Specifies the unnumbered interface name of the neighbor.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.

VPNV4 address-family, peer group, and BFD are not supported for BGP unnumbered.

Examples

Enabling and removing the address-family capability and exchange of information on 1/1/1.

```
switch(config-bgp)# neighbor 1/1/1 remote-as 1
switch(config-bgp-ipv4-uc)# neighbor 1/1/1 activate
switch(config-bgp-ipv4-uc)# no neighbor 1/1/1 activate
```

Enabling and removing the address-family capability and exchange of information on 1/1/1.1.

```
switch(config-bgp)# neighbor 1/1/1.1
switch(config-bgp-ipv6-uc)# neighbor 1/1/1.1 activate
switch(config-bgp-ipv6-uc)# no neighbor 1/1/1.1 activate
```

Enabling and removing the address-family capability and exchange of information on 1/1/1:1.

```
switch(config-bgp)# neighbor 1/1/1:1
switch(config-bgp-ipv4-uc)# neighbor 1/1/1:1 activate
switch(config-bgp-ipv4-uc)# no neighbor 1/1/1:1 activate
```

Enabling and removing the address-family capability and exchange of information on 1.1.1.1.

```
switch(config-bgp-l2vpn-evpn)# neighbor 1.1.1.1 activate
switch(config-bgp-l2vpn-evpn)# no neighbor 1.1.1.1 activate
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority

6300, except for S3L75A, S3L76A, S3L77A
6400
8100
8325
8325H
8325P
8360
9300
9300S
10000
	`config-bgp-ipv4-uc`
	`config-bgp-ipv6-uc`
	`config-bgp-l2vpn-evpn`
		Administrators or local user group members with execution rights for this command.

neighbor advertisement-interval

```
neighbor {<IP-ADDRESS>|<IFNAME>|<PEER-GROUP-NAME>} advertisement-interval <INTERVAL>
no neighbor {<IP-ADDRESS>|<IFNAME>|<PEER-GROUP-NAME>} advertisement-interval
```

Description

Sets the advertisement interval, which defines the length of time between transmission of BGP routing updates.

The **no** form of this command restores the default value. Default values are 30 seconds for external BGP peer and 5 seconds for internal BGP peer.

Parameter	Description
<IP-ADDR>	Specifies an IP address in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255, or IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a Peer-Group. This is not supported on BGP unnumbered.
<INTERVAL>	Specifies the advertisement interval in seconds. Range: 0 to 600. Default: 30 for external BGP peer and 5 for internal BGP peer.

Examples

```
switch(config-bgp-ipv4-uc)# neighbor 1.1.1.1 advertisement-interval 20
switch(config-bgp-ipv4-uc)# no neighbor 1.1.1.1 advertisement-interval
```

```
switch(config-bgp-ipv4-uc)# neighbor pg advertisement-interval 50
switch(config-bgp-ipv4-uc)# no neighbor pg advertisement-interval
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp-ipv4-uc` `config-bgp-ipv6-uc`	Administrators or local user group members with execution rights for this command.

neighbor add-paths

```
neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} add-paths {send | recv | both}
no neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} add-paths {send | recv | both}
```

Description

Enables additional path capability for BGP as described in RFC 7911. This allows BGP peer to send, receive, or send and receive multiple paths for the same address prefix without the subsequent advertisements implicitly replacing any previous paths. The additional path includes the first (N-1) best paths, which means that the total paths for an address prefix received by a BGP speaker will include the best path and the additional paths determined by its BGP peer.

With additional path feature, each path is identified by a path identifier in addition to the address prefix. To use this command, the backup path of BGP next-hop must be different than the primary path.

The **no** form of this command disables the additional path feature.

Parameter	Description
<IP-ADDR>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.
`add-paths {send | recv | both}`	Configures the additional paths in one of the following ways: `send`—Enables the neighbor to send the additional paths. `recv`—Enables the neighbor to receive the additional paths. `both`—Enables the neighbor to send and receive the additional paths.

Examples

Enabling BGP neighbor to send the additional paths:

```
switch(config)# router bgp 100
switch(config-bgp)# address-family ipv4 unicast
switch(config-bgp-ipv4-uc)# neighbor 1.1.1.1 add-paths send
```

Disabling BGP neighbor to send the additional paths:

```
switch(config)# router bgp 100
switch(config-bgp)# address-family ipv4 unicast
switch(config-bgp-ipv4-uc)# no neighbor 1.1.1.1 add-paths send
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp-ipv4-uc` `config-bgp-ipv6-uc`	Administrators or local user group members with execution rights for this command.

neighbor add-paths advertise-best

```
neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} add-paths advertise-best <2-4>
no neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} add-paths advertise-best <2-4>
```

Description

Controls the number of best BGP Paths to be advertised by a BGP speaker to a BGP peer. When enabled, it allows BGP speaker to advertise more than one best paths for the same address prefix. The total paths for an address prefix will include the best path and the additional paths.

The **no** form of this command removes the advertise best path configuration.

Parameter	Description
<IP-ADDR>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.
`advertise-best` *<2-4>*	Specifies the number of best BGP paths to be advertised to a BGP Peer. Range: 2 to 4. Default: 2.

Examples

Setting the number of best paths to send to the neighbor:

```
switch(config)# router bgp 100
switch(config-bgp)# address-family ipv4 unicast
switch(config-bgp-ipv4-uc)# neighbor 1.1.1.1 add-paths advertise-best 3
```

Removing the advertise best path configuration:

```
switch(config)# router bgp 100
switch(config-bgp)# address-family ipv4 unicast
switch(config-bgp-ipv4-uc)# no neighbor 1.1.1.1 add-paths advertise-best 3
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp-ipv4-uc` `config-bgp-ipv6-uc`	Administrators or local user group members with execution rights for this command.

neighbor allowas-in

`neighbor {`*<IP-ADDRESS>*`|`*<IFNAME>*`|`*<LIMIT>*

```
no neighbor {<IP-ADDRESS>|<IFNAME>|<LIMIT>
```

Description

Specifies the number of times that the AS path of a received route can contain the AS number of the recipient BGP speaker and still be accepted. When this configuration is applied to a peer-group, all the neighbors that are part of the peer-group inherit this setting.

The **no** form of this command restores the default setting, which is to reject as a loop any route where the path contains the speaker AS number.

Parameter	Description
<IP-ADDRESS>	Specifies the neighbor IP address in the IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255, or in the IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F. **NOTE:** IPv6 MP-BGP peering must not be used for L2VPN EVPN address family, because VXLAN tunnel interface does not support IPv6 addresses.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group.
<LIMIT>	Specifies the number of times that the AS path of a received route can contain the AS number of the recipient BGP. Range: 1 to 10.

Examples

```
switch(config-bgp-ipv4-uc)# neighbor 1.1.1.1 allowas-in 5
switch(config-bgp-ipv4-uc)# no neighbor 1.1.1.1 allowas-in
```

```
switch(config-bgp-ipv6-uc)# neighbor 2001:0db8:85a3::8a2e:0370:7334 allowas-in 5
switch(config-bgp-ipv6-uc)# no neighbor 2001:0db8:85a3::8a2e:0370:7334 allowas-in
```

```
switch(config-bgp-ipv4-uc)# neighbor PG allowas-in 5
switch(config-bgp-ipv4-uc)# no neighbor PG allowas-in
```

```
switch(config-bgp-l2vpn-evpn)# neighbor 1.1.1.1 allowas-in 5
switch(config-bgp-l2vpn-evpn)# no neighbor 1.1.1.1 allowas-in
```

```
switch(config-bgp-l2vpn-evpn)# neighbor PG allowas-in 5
switch(config-bgp-l2vpn-evpn)# no neighbor PG allowas-in
```

```
switch(config-bgp-l2vpn-evpn)# neighbor 2001:0db8:85a3::8a2e:0370:7334 allowas-in
5
switch(config-bgp-l2vpn-evpn)# no neighbor 2001:0db8:85a3::8a2e:0370:7334 allowas-
in
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority

6300, except for S3L75A, S3L76A, S3L77A
6400
8100
8325
8325H
8325P
8360
9300
9300S
10000
	`config-bgp-ipv4-uc`
	`config-bgp-ipv6-uc`
	`config-bgp-l2vpn-evpn`
		Administrators or local user group members with execution rights for this command.

neighbor ao

```
neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} ao <keychain-name>
   accept-ao-mismatch
   include-tcp-options
   no ...
```

Description

Enables TCP Authentication Option (TCP-AO) authentication on a TCP connection between two BGP neighbors. To disable this function, use the **no** form of this command.

Parameter	Description
<IP-ADDRESS>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.
<keychain-name>	Name of the keychain for the neighbor. The keychain allows keys to be configured with different valid lifetimes. This mechanism provides a way for a set of keys to be rotated and hence protect against long-lived-key attacks. At any given time only one key is selected as active-key and keys are valid for a duration of the defined send-lifetime. If the send-lifetime and the accept-lifetimes are not configured for the key, the key is considered to be valid for infinite lifetime. When multiple keys are configured, its recommended that keys

Parameter	Description
	overlap in their send-lifetimes so that the key rollover occurs at the start of the next key's send-lifetime. This allows for a continuous key usage by TCP-AO.
`accept-ao-mismatch`	Accept incoming TCP segments without TCP-AO option. If enabled, the device will accept a connection from the peer *even if the received TCP packets do not contain the TCP-AO option*.
`include-tcp-options`	Include the TCP header options for MAC calculation. Note that enabling this setting will immediately reset the neighbor session. This setting is disabled by default.
`no ...`	Negates any configured parameter.

Usage

TCP-AO authentication can not be used with the neighbor password feature. When TCP-AO is applied to a peer-group, all the neighbors in peer-group will inherit the peer-group configuration unless there is a configuration specific to an individual neighbor. If a peer-group is configured with the neighbor password feature but the neighbors that belong to that peer-group are configured with TCP-AO, the TCP-AO configuration will be rejected. Similarly, If a peer-group is configured to use TCP-AO authentication, the neighbors that belong to that peer-group will reject the neighbor password.

The TCP-AO feature takes a keychain as a parameter. The key will not be valid until a Recv-D, Send-ID, and send lifetime is configured. The supported cryptographic algorithms for TCP-AO are:

- HMAC-SHA-1-96 based on [RFC2104] and [FIPS-180-3]
- AES-128-CMAC-96 based on [NIST-SP800-38B][FIPS197]

Examples

```
switch(config)# keychain bgpkeys
switch(config-keychain)# key 1
switch(config-keychain-key)# send-lifetime start-time 10:10:10 10/25/2022 duration
infinite
switch(config-keychain-key)# accept-lifetime start-time 10:10:10 10/25/2022
duration infinite
switch(config-keychain-key)# send-id 10
switch(config-keychain-key)# recv-id 10
switch(config-keychain-key)# cryptographic-algorithm aes-cmac-128
switch(config-keychain-key)# key-string plaintext qwer
switch(config)# router bgp 1
switch(config-bgp)# neighbor 1.1.1.1 ao bgpkeys
switch(config-bgp)# no neighbor 1.1.1.1 ao accept-ao-mismatch
switch(config-bgp)# no neighbor 1.1.1.1 ao include-tcp-option
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.11 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

neighbor capability orf prefix-list

```
neighbor {<IP-ADDRESS>|<IFNAME>} capability orf prefix-list {send | receive | both}
no neighbor {<IP-ADDRESS>|<IFNAME>} capability orf prefix-list {send | receive | both}
```

Description

Enables the Outbound Route filtering (ORF) capability with the neighbor in one of the three available modes. The available modes are send, receive, and both. The ORF capability is executed based on prefix list only.

The Outbound Route Filtering (ORF) capability provides a mechanism for a BGP speaker to send a set of Outbound Route Filters (ORFs) that can be used by its BGP peer to filter its outbound routing updates to the speaker. This is a filtering method used to reduce the computation on the router receiving the route.

The **no** form of this command disables the ORF capability.

Parameter	Description
<IP-ADDRESS>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
`capability orf prefix-list {send | receive | both}`	Enables ORF prefix list capability with the neighbor in one of the following modes: ■ **send** - Enables the ORF prefix list capability in send mode. ■ **receive** - Enables the ORF prefix list capability in receive mode. ■ **both**- Enables the ORF prefix list capability in both send and receive mode.

Examples

Enabling the ORF prefix list capability in both send and receive mode:

```
switch(config-bgp-ipv4-uc)# neighbor 1.1.1.1 capability orf prefix-list both
```

Enabling the ORF prefix list capability in send mode:

```
switch(config-bgp-ipv4-uc)# neighbor 1.1.1.1 capability orf prefix-list send
```

Disabling the ORF prefix list capability in both send and receive mode: :

```
switch(config-bgp-ipv4-uc)# no neighbor 1.1.1.1 capability orf prefix-list both
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp-ipv4-uc` `config-bgp-ipv6-uc`	Administrators or local user group members with execution rights for this command.

neighbor default-originate

`neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} default-originate [route-map <MAP-NAME>]`
`no neighbor {<IP-ADDRESS>|<IFNAME>|<PEER-GROUP-NAME>} default-originate [route-map <MAP-NAME>]`

Description

Enables the local router to send the default route 0.0.0.0 to a neighbor. The neighbor can then use this route to reach the router when all other routes are unavailable. Use the `route-map` option to configure the route map to modify the default route attributes.

The **no** form of this command disables this feature.

Parameter	Description
<IP-ADDR>	Specifies an IP address in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255, or IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F.

Parameter	Description
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.
<MAP-NAME>	Sets the route map to modify the default route attributes.

Examples

```
switch(config-bgp-ipv4-uc)# neighbor 1.1.1.1 default-originate
switch(config-bgp-ipv4-uc)# no neighbor 1.1.1.1 default-originate
```

```
switch(config-bgp-ipv4-uc)# neighbor PG default-originate
switch(config-bgp-ipv4-uc)# no neighbor PG default-originate
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp-ipv4-uc` `config-bgp-ipv6-uc`	Administrators or local user group members with execution rights for this command.

neighbor ebgp-multihop

```
neighbor {<IP-ADDR> | <PEER-GROUP-NAME>} ebgp-multihop <ttl>
no neighbor {<IP-ADDR> | <PEER-GROUP-NAME>} ebgp-multihop <ttl>
```

Description

Enables BGP to establish a session with external peers residing on networks that are not directly connected. By default, BGP can only establish sessions with external BGP peers that are directly connected. If external BGP peer are not connected directly, then ebgp-multihop is enabled to initiate

the connection with that external peer. If ebgp-multihop is disabled and external BGP peers are indirectly connected, then BGP peer session is not established.

The **no** form of this command disables the peer ebgp-multihop feature.

Parameter	Description
<IP-ADDR>	Specifies an IP address.
<PEER-GROUP-NAME>	Specifies a peer group.
<ttl>	Specifies the ttl value.

Examples

Enabling BGP to establish connection with external peers residing on networks that are not directly connected:

```
switch(config-bgp)# neighbor 1.1.1.1 ebgp-multihop 5
```

Disabling BGP to establish connection with external peers residing on networks that are not directly connected:

```
switch(config-bgp)# no neighbor 1.1.1.1 ebgp-multihop
```

Enabling BGP to establish connection with external peers residing on networks that are not directly connected:

```
switch(config-bgp)# neighbor pg ebgp-multihop 5
```

Disabling BGP to establish connection with external peers residing on networks that are not directly connected:

```
switch(config-bgp)# no neighbor pg ebgp-multihop 5
```

Command History

Release	Modification
10.16	Syntax updated to use <*ttl*> instead of <*HOP-COUNT*>.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A,	`config-bgp`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000		

neighbor fall-over

```
neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} fall-over
no neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} fall-over
```

Description

Enables BGP fast peering session deactivation. When neighbor fall-over is configured, the BGP process monitors the RIB and if the route to peer is not present in the routing table, it immediately deactivates the peer session without waiting for the hold down timer. It is disabled by default.

The **no** form of this command disables this feature.

Parameter	Description
<IP-ADDR>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.

Usage

Neighbor fall-over does not track connected or static routes to peers. However, this is not an issue when IBGP peering is using a loopback interface. To force a fall-over for connected and static routes, use the command `neighbor fall-over bfd`.

Examples

```
switch(config-bgp)# neighbor 1.1.1.1 fall-over
switch(config-bgp)# no neighbor 1.1.1.1 fall-over
```

```
switch(config-bgp)# neighbor PG fall-over
switch(config-bgp)# no neighbor PG fall-over
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

neighbor fall-over bfd

```
neighbor {<IPV4-ADDR>|<IPV6-ADDR>|<PEER-GROUP-NAME>} fall-over bfd
no neighbor {<IPV4-ADDR>|<IPV6-ADDR>|<PEER-GROUP-NAME>} fall-over bfd
```

Description

Enables BGP to register with BFD to receive fast peering session deactivation messages from BFD. You can either configure BFD support for BGP per neighbor or peer-group.

The **no** form of this command disables BGP for BFD.

Multihop BFD is not supported for BGP.

Parameter	Description
<IPV4-ADDR>	Specifies the neighbor address in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.
<IPV6-ADDR>	Specifies the neighbor address in IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F. This parameter applies only to HPE Aruba Networking 6300 and 6400 Switch series.
<PEER-GROUP-NAME>	Specifies a peer group.

Examples

Enabling BFD for BGP neighbor with IPv4 address:

```
switch(config-bgp)# neighbor 1.1.1.1 fall-over bfd
```

Enabling BFD for BGP neighbor with IPv6 address (applies only to HPE Aruba Networking 6300 and 6400 Switch series):

```
switch(config-bgp)# neighbor 1000::1 fall-over bfd
```

Enabling BFD for peer group:

```
switch(config-bgp)# neighbor PG fall-over bfd
```

Disabling BFD for BGP per neighbor IPv4 address:

```
switch(config-bgp)# neighbor 1.1.1.1 fall-over bfd
```

Disabling BFD for BGP per neighbor with IPv6 address (applies only to HPE Aruba Networking 6300 and 6400 Switch series):

```
switch(config-bgp)# no neighbor 1000::1 fall-over bfd
```

Disabling BFD for peer group:

```
switch(config-bgp)# no neighbor PG fall-over bfd
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-bgp	Administrators or local user group members with execution rights for this command.

neighbor graceful-shutdown

```
neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} graceful-shutdown
     [local-preference <LOCAL-PREF>
     | <CONFIG-DELAY>
    |<LOCAL-PREF> ]
no neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} graceful-shutdown
     [local-preference <LOCAL-PREF>
     | <CONFIG-DELAY>
     |<LOCAL-PREF> ]
```

Description

Configures the wait time before shutting down the BGP neighbor session, and can also configure the local preference value to be advertised before graceful shutdown.

The **no** form of this command sets the wait time to the default value of 180 seconds and the local-preference value to the default of 0.

Parameter	Description
<IP-ADDR>	Specifies an IP address in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255, or IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.
`local-preference`*<LOCAL-PREF>*	Specifies the local preference value for exporting the iBGP routes during a graceful shutdown. The lower the value, the lower the local preference. A value of 0 indicates that the route is the least preferred. Range: 0 to 4294967295. Default: 0.
<CONFIG-DELAY>	Specifies the time to wait before shutting down the neighbor in seconds. Range: 10 to 1200. Default: 180.

Usage

If the graceful shutdown timer has already started and the administrator configures a command that triggers a session restart, traffic loss can occur if the graceful shutdown delay is not sufficient for the BGP peers to converge to a new route.

On each Autonomous System Boundary Router (ASBR) supporting the graceful shutdown receiver procedure, an inbound BGP route policy must be applied on all EBGP sessions of the ASBR.

The policy must match the GSHUT community and lower the precedence of the route by changing the route attributes.

The Graceful-Shutdown feature does not work for reflected routes because the route reflector (RR) does not modify local-preference attribute. The routes, originated by the RR, carry the GSHUT local-preference value. As per the RFC 4456, when an RR reflects a route, it should not modify the following path attributes:

- NEXT-HOP
- AS-PATH

- LOCAL-PREF
- MED

Their modification could potentially result in routing loops. In this situation, apply on the RR an inbound BGP route policy, meeting the following conditions:

- Match the graceful-shutdown community.
- Set the local preference attributes of the paths tagged with the graceful-shutdown community to a lower value than other routes to the same destination.

Examples

Setting the wait time delay:

```
switch(config-bgp)# neighbor 1.1.1.1 graceful-shutdown 10
```

Setting the local-preference value:

```
switch(config-bgp)# neighbor 1.1.1.1 graceful-shutdown local-preference 100
```

Setting the wait time delay and local-preference value:

```
switch(config-bgp)# neighbor 1.1.1.1 graceful-shutdown 10 local-preference 100
```

Setting the wait time delay to the default of 180 seconds:

```
switch(config-bgp)# no neighbor 1.1.1.1 graceful-shutdown 10
```

Setting the local-preference value to default of 0:

```
switch(config-bgp)# no neighbor 1.1.1.1 graceful-shutdown local-preference 100
```

Setting the wait time delay and local-preference value to defaults:

```
switch(config-bgp)# no neighbor 1.1.1.1 graceful-shutdown 10 local-preference 100
```

Complete deletion:

```
switch(config-bgp)# no neighbor 1.1.1.1 graceful-shutdown
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

neighbor interface

```
neighbor {<IP-ADDR>|<IFNAME>} interface <ifname>
[no] neighbor {<IP-ADDR>|<IFNAME>} interface <ifname>
```

Description

Configure a BGP neighbor by specifying the interface name on which the neighbor is connected instead of the neighbor's IP address. When an unnumbered interface is specified to the **neighbor** command, a static route is auto-provisioned for the neighbor with the unnumbered interface as a destination. A static route is also auto-provisioned if the neighbor IP address is learned via LLDP.

If OSPF is enabled on an unnumbered interface then the auto-provisioning is not supported as the route to the BGP peer is expected to be learned via OSPF.

Peer-group is not supported for BGP peering on the IP unnumbered interface.

The **no** form of this command disables this feature.

Parameter	Description
<IP-ADDR>	Specifies an IP address.

Examples

Configuring unnumbered BGP peering on interface 1/1/1:

```
switch(config)# router bgp 100
switch(config-bgp)# neighbor 1.1.1.1 remote-as 100
switch(config-bgp)# neighbor 1.1.1.1 interface 1/1/1
```

Disabling this feature:

```
switch(config)# router bgp 100
switch(config-bgp)# no neighbor 1.1.1.1 interface 1/1/1
```

Command History

Release	Modification
10.14.1000	Command introduced

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

neighbor listen ip-range

```
neighbor <PEER-GROUP-NAME> listen ip-range <IP-ADDR>|<MASK> {as-range <AS-RANGE>|limit <LIMIT>}
no neighbor <PEER-GROUP-NAME> listen ip-range <IP-ADDR>|<MASK> {as-range <AS-RANGE>|limit <LIMIT>}
```

Description

Configures BGP dynamic neighbors as ranges of remote addresses with associated peer groups.

The **no** form of this command restores default behavior.

- Each range of remote addresses is configured as a remote address prefix.
- Any BGP peer with a remote address that matches the remote address prefix becomes a member of the associated peer group.

Parameter	Description
<PEER-GROUP-NAME>	Specifies peer group. This is not supported on BGP unnumbered.
<IP-ADDR>|*<MASK>*	Specifies subnet range.
<AS-RANGE>	Specifies AS number as a range in integer or dotted format.
<LIMIT>	Specifies maximum number of peers. Range: 1 to 256.

Restrictions

- Dynamic peers are always passive. Outbound connections to dynamic peers are not supported.
- Dynamic BGP peering is only compatible with peer-groups

- Disabling partial AS range is not supported. The exact value that is configured must be used.
  - When disabling AS range, CLI must use the same AS range that was used when first configured. For example, if AS range "1-4" is configured, when disabling, "1-4" must be used ("1,2,3,4" is not supported).

Configuring overlapping peer ranges with different remote address prefix lengths is not recommended. Peer range configuration is recommended when peer ranges do not overlap.

Usage

- All supported address-families are activated on a dynamic peer for negotiation by default.
- If an incoming connection matches multiple peer range entries, the entry with the longest remote address prefix is selected.
- AS ranges are used to match remote AS presented by connecting peers. Remote AS matching with ASes or AS ranges in this list will be accepted.
  - AS range only applies to dynamic peers.
- The limit option is used to set the maximum number of dynamic BGP peers within the peer range. The default is 512 if no limit is set.
  - If the limit is reached, BGP rejects incoming connections from new dynamic BGP peers until BGP session termination causes the number of dynamic BGP peers to fall below the limit.
  - If the limit is reduced below the current number of dynamic BGP peers, BGP will reject incoming connections from new dynamic BGP peers until the number of dynamic BGP peers falls below the new limit. BGP will not terminate existing BGP sessions with dynamic BGP peers in this case. If an existing BGP session gets terminated, that session will not re-establish until the number of BGP sessions falls below the limit.
- After dynamic peer is configured, additional configuration is required on the peer-group as a whole. Individual member groups are incompatible. For example, the neighbor shutdown command can be executed on a peer-group, but not on individual members of the peer-group.
- When a peer is configured as dynamic and is in an established state, a shutdown is required before reconfiguring as static.
- Connect-retry interval is recommended to be configured with a smaller value than the default value on the active peer.
- When a set of valid and invalid AS values are issued (separated by commas), only the valid values are accepted.
- When the AS range parameter is not explicitly configured in dynamic bgp peering, iBGP session comes up, eBGP session does not. If there are no configured remote AS or AS list entries, DC-BGP assumes that any peer is an iBGP peer.

Examples

```
switch(config-bgp)# neighbor pg listen ip-range 192.168.0.0/16
```

```
switch(config-bgp)# no neighbor pg listen ip-range 192.168.0.0/16
```

Command History

Release	Modification
10.11	Command introduced

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

neighbor local-as

```
neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} local-as <AS-NUMBER> [no-prepend[replace-as[dual-as]]]
no neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} local-as [no-prepend[replace-as[dual-as]]]
```

Description

Configures an alternate local AS number that can be used to establish a session with a peer, allowing a router to appear to be a member of a second autonomous system (AS), and its real AS.

Local AS allows two autonomous systems to merge without modifying peering arrangements. This command is valid only for external peers.

The **no** form of this command restores the default, which is for a peering session to be established using the primary AS (primary AS is the AS number specified at the time of neighbor creation using the command `neighbor remote-as`).

Parameter	Description
<IP-ADDR>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.
`local-as` *<AS-NUMBER>*	Specifies a 4-byte AS number in asplain format (z), or asdot format (x.y), where z is a number from 1 to 4294967295 and x and y are 16-bit numbers.
`no-prepend`	Optional. When you include this parameter, the switch will not prepend local ASN value to the AS path list for all the routes received from the peer.

Parameter	Description
replace-as	Optional. Replaces the configured router AS with the local AS in the EBGP updates.
dual-as	Optional. Use this option to form the neighborship with any of the two AS numbers; the configured router ASN or the local ASN.

Examples

Configuring an alternate local AS number that can be used to establish a session with a peer:

```
switch(config-bgp)# neighbor 1.1.1.1 local-as 200
```

Removing the alternate local AS number that can be used to establish a session with a peer:

```
switch(config-bgp)# no neighbor 1.1.1.1 local-as 200
```

Configuring an alternate local AS number that can be used to establish a session with a peer:

```
switch(config-bgp)# neighbor pg local-as 200 no-prepend
```

Removing the alternate local AS number that can be used to establish a session with a peer:

```
switch(config-bgp)# no neighbor pg local-as 200 no-prepend
```

Configuring an alternate local AS number that can be used to establish a session with a peer:

```
switch(config-bgp)# neighbor pg local-as 200 no-prepend replace-as
```

Removing the alternate local AS number that can be used to establish a session with a peer:

```
switch(config-bgp)# no neighbor pg local-as 200 no-prepend replace-as
```

Configuring an alternate local AS number that can be used to establish a session with a peer:

```
switch(config-bgp)# neighbor pg local-as 200 no-prepend replace-as dual-as
```

Removing the alternate local AS number that can be used to establish a session with a peer:

```
switch(config-bgp)# no neighbor pg local-as 200 no-prepend replace-as dual-as
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

neighbor maximum-prefix

```
neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>}
        maximum-prefix <MAXIMUM> [threshold <THRESHOLD>]
        [restart <INTERVAL>] [warning-only]
no neighbor  {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} maximum-prefix<MAXIMUM> [threshold
<THRESHOLD>]
        [restart <INTERVAL>] [warning-only]
```

Description

Sets the maximum number of prefixes that can be received from a neighbor.

The session is re-established only if the number of routes received from the BGP peer does not exceed the configured prefix limit. When the restart timer is configured, sessions are automatically re-established when the timer expires. By default, the device accepts 128,000 prefixes from a BGP neighbor with a threshold value of 75%. A warning message is generated when the number of prefixes per neighbor reaches 75% of default prefix limit. Another warning message is generated when the default prefix limit is reached. A maximum prefix limit applies at the neighbor level. If maximum-prefix limit exceeds any configured limit, the BGP session with that neighbor will be brought down.

The **no** form of this command disables the maximum number of prefixes limit.

Parameter	Description
<IP-ADDR>	Specifies the IP address of the neighbor in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255, or IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a Peer-Group. This is not supported on BGP

Parameter	Description
	unnumbered.
`maximum-prefix <MAXIMUM>`	Specifies the maximum number of prefixes allowed from the specified neighbor. Range: 1 to 128000. Default: 128000.
`threshold <THRESHOLD>`	Specifies at what percentage of `MAXIMUM` a warning message is generated. Range: 1 to 100. Default: 75. For example, if `MAXIMUM` is set to 1000 and threshold is 70, the router generates a warning message when the number of BGP learned routes from the neighbor exceeds 70 percent of 1000 (700) routes.
`restart <INTERVAL>`	Specifies interval in seconds for restarting the BGP connection when the prefix limit is exceeded. Range: 30 to 65535.
`warning-only`	Specifies generating and logging a warning message without disconnecting the BGP session when the prefix limit is exceeded.

Examples

Setting the prefix limit to 1000 prefixes:

```
switch(config-bgp-ipv4-uc)# neighbor 10.0.0.1 maximum-prefix 1000
```

Enabling logging of a warning message when more than 1000 prefixes are received:

```
switch(config-bgp-ipv4-uc)# neighbor 10.0.0.1 maximum-prefix 1000 warning-only
```

Setting the prefix limit to 1000 prefixes and enabling logging of a warning message when 500 prefixes are received:

```
switch(config-bgp-ipv4-uc)# neighbor 10.0.0.1 maximum-prefix 1000 threshold 50
```

Setting the prefix limit to 1000 prefixes and enabling logging of a warning message when 500 prefixes are received and a second warning when the prefix limit is exceeded without disconnecting the session:

```
switch(config-bgp-ipv4-uc)# neighbor 10.0.0.1 maximum-prefix 1000 threshold 50
warning-only
```

Removing the threshold value:

```
switch(config-bgp-ipv4-uc)# no neighbor 10.0.0.1 maximum-prefix 1000 threshold 50
```

Disabling the maximum-prefix feature:

```
switch(config-bgp-ipv4-uc)# no neighbor 10.0.0.1 maximum-prefix
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp-ipv4-uc` `config-bgp-ipv6-uc`	Administrators or local user group members with execution rights for this command.

neighbor next-hop-self

```
neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>}next-hop-self
no neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} next-hop-self
```

Description

Configures the router as the next hop for a BGP-speaking neighbor or peer group, and enables BGP to send itself as the next hop for advertised routes.

The **no** form of this command resets the peer next-hop-self status to default. The next hop is generated based on the IP.

Parameter	Description
<IP-ADDR>	Specifies the neighbor's IP address in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255, or IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.
`all`	Applies the configuration to all route-reflector clients.

Usage

- An administrator uses this command to make a BGP speaker fill its address when advertising routes to a BGP peer.
- This command is useful in non-meshed networks where BGP neighbors may not have direct access to all other neighbors on the same IP subnet.

- While advertising routes to eBGP peers, the next-hop is set to self IP by default. The default behavior can be changed by configuring next-hop-unchanged.
- While advertising routes to iBGP peers, the next-hop is kept unchanged by default. The default behavior can be changed by configuring next-hop-self.

BGP L2VPN EVPN context:

- The next-hop-self command under l2vpn-evpn address-family, in addition to replacing the next-hop IP of EVPN routes with self IP, replaces the routers-mac extended community of EVPN routes with self MAC.

Examples

BGP IPv4 UC context:

Setting and resetting the router as the next hop self for neighbor 1.1.1.1:

```
switch(config-bgp-ipv4-uc)# neighbor 1.1.1.1 next-hop-self
switch(config-bgp-ipv4-uc)# no neighbor 1.1.1.1 next-hop-self
```

Setting and resetting the router as the next hop self for its peer group:

```
switch(config-bgp-ipv4-uc)# neighbor pg next-hop-self
switch(config-bgp-ipv4-uc)# no neighbor pg next-hop-self
```

BGP L2VPN EVPN context:

Supported only on the HPE Aruba Networking 8325 Switch Series

Setting and resetting the router as the next hop self for neighbor 1.1.1.1:

```
switch(config-bgp-l2vpn-evpn)# neighbor 1.1.1.1 next-hop-self
switch(config-bgp-l2vpn-evpn)# no neighbor 1.1.1.1 next-hop-self
```

Setting and resetting the router as the next hop self for its peer group:

```
switch(config-bgp-l2vpn-evpn)# neighbor pg next-hop-self
switch(config-bgp-l2vpn-evpn)# no neighbor pg next-hop-self
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.9	Added support for BGP L2VPN EVPN address family.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority

6300, except for S3L75A, S3L76A, S3L77A
6400
8100
8325
8325H
8325P
8360
9300
9300S
10000
	`config-bgp-ipv4-uc`
	`config-bgp-ipv6-uc`
	`config-bgp-l2vpn-evpn`
		Administrators or local user group members with execution rights for this command.

neighbor next-hop-unchanged

```
neighbor {<IP-ADDRESS>|<IFNAME>} next-hop-unchanged
no neighbor  {<IP-ADDRESS>|<IFNAME>} next-hop-unchanged
```

Description

Enables the neighbor to preserve next-hop while advertising routes to eBGP peers, in the L2VPN EVPN address-family.

The **no** form of this command resets the peer next-hop-unchanged status to default.

Parameter	Description
<IP-ADDRESS>	Specifies the neighbor IP address in the IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255, or in the IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F. **NOTE:** IPv6 MP-BGP peering must not be used for L2VPN EVPN address family, because VXLAN tunnel interface does not support IPv6 addresses.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.

Examples

```
switch(config-bgp-l2vpn-evpn)# neighbor 1.1.1.1 next-hop-unchanged
```

```
switch(config-bgp-l2vpn-evpn)# neighbor 2001:0db8:85a3::8a2e:0370:7334 next-hop-
unchanged
```

The **next-hop-unchanged** option needs to be explicitly configured to preserve the nexthop while advertising routes to eBGP peers, in the L2VPN EVPN address-family. For example:

```
   router bgp 1
neighbor 1.1.1.1 remote-as 2
      address-family l2vpn evpn
```

```
         neighbor 1.1.1.1 activate
neighbor 1.1.1.1 next-hop-unchanged
         neighbor 1.1.1.1 send-community extended
      exit-address-family
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp-l2vpn-evpn`	Administrators or local user group members with execution rights for this command.

neighbor orf prefix-list in

`neighbor {<IP-ADDRESS>|<IFNAME>} orf-prefix-list <PREFIX-LIST-NAME> in`
`no neighbor {<IP-ADDRESS>|<IFNAME>} orf-prefix-list <PREFIX-LIST-NAME> in`

Description

Applies an inbound prefix list filter to filter the distribution of BGP neighbor information.

The **no** form of this command restores the default behavior of not applying the prefix list filter.

This command must be used only along with the ORF capability to take effect.

Parameter	Description
<IP-ADDRESS>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
`orf-prefix-list` *PREFIX-LIST-NAME>*	Sends the prefix list name to be filtered.

Usage

To use this command, the following conditions must be met:

- If route-map inbound is also applied on multiple neighbors along with ORF, then the route-map name must be common on all the neighbors.
- If route-map inbound is also applied on an IPv6 AF BGP neighbor, then the route-map sequence number with value 1 cannot be used.

Examples

Applying the inbound prefix list filter:

```
switch(config-bgp-ipv4-uc)# neighbor 1.1.1.1 orf-prefix-list ABC in
```

Removing the inbound prefix list filter:

```
switch(config-bgp-ipv4-uc)# no neighbor 1.1.1.1 orf-prefix-list ABC in
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp-ipv4-uc` `config-bgp-ipv6-uc`	Administrators or local user group members with execution rights for this command.

neighbor passive

```
neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>}  passive
no neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} passive
```

Description

Configures a specific neighbor, or all neighbors in a peer-group, as passive, which means that they will not initiate the TCP session.

The **no** form of this command enables the neighbor to initiate the TCP session.

Parameter	Description
<IP-ADDRESS>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group.

Examples

Configure a specific neighbor as passive:

```
switch(config-bgp)# neighbor 1.1.1.1 passive
```

Enabling the local router to initiate the TCP session for a specified neighbor:

```
switch(config-bgp)# no neighbor 1.1.1.1 passive
```

Configure a specific peer group as passive:

```
switch(config-bgp)# neighbor pg passive
```

Enabling the local router to initiate the TCP session for a specified peer group:

```
switch(config-bgp)# no neighbor pg passive
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

neighbor password

```
neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>}      password [{ciphertext | plaintext}
<PASSWORD>]
no neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>}
            password [ciphertext <PASSWORD>]
```

Description

Enables message digest5 (MD5) authentication on a TCP connection between two BGP neighbors. When the password is applied to a peer-group, all the neighbors that are part of peer-group inherit the configured setting.

The **no** form of this command removes the neighbor password.

Parameter	Description
<IP-ADDR>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a Peer-Group. This is not supported on BGP unnumbered.
{ciphertext | plaintext}	Selects the password format.
<PASSWORD>	Specifies the password.

When the password is not provided on the command line, plaintext password prompting occurs upon pressing Enter. The entered password characters are masked with asterisks.

Examples

Enabling message digest5 (MD5) authentication for a neighbor with a provided plaintext password:

```
switch(config-bgp)# neighbor 1.1.1.1 password plaintext doubt_Plane#93
```

Enabling message digest5 (MD5) authentication for a neighbor with a prompted plaintext password:

```
switch(config-bgp)# neighbor 1.1.1.5 password
Enter the neighbor password: *************
Re-Enter the neighbor password: *************
```

Enabling message digest5 (MD5) authentication for a peer group with a provided plaintext password:

```
switch(config-bgp)# neighbor pg_3 password plaintext doubt_Plane#93
```

Disabling message digest5 (MD5) authentication for a neighbor:

```
switch(config-bgp)# no neighbor 1.1.1.5 password
```

Disabling message digest5 (MD5) authentication for a peer group:

```
switch(config-bgp)# no neighbor pg_3 password
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-bgp	Administrators or local user group members with execution rights for this command.

neighbor port

```
neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} port <NUMBER>
no neighbor{<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} port <NUMBER>
```

Description

Sets a custom TCP port on which to communicate with the BGP neighbor.

When this configuration is applied for peer-group, all the neighbors that are part of peer-group will inherit this setting. Though the neighbor inherits the configuration from the peer-group, the neighbor-specific command, if configured, takes precedence.

This setting only takes effect after a hard reset of the session.

The **no** form of this command allows a random TCP port to be selected for the communication with the BGP neighbor.

Parameter	Description
<IP-ADDRESS>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.
port *<NUMBER>*	Specifies a TCP port number. Range: 0 to 65535.

Examples

```
switch(config-bgp)# neighbor 1.1.1.1 port 1500
switch(config-bgp)# no neighbor 1.1.1.1 port
```

```
switch(config-bgp)# neighbor PG port 1500
switch(config-bgp)# no neighbor PG port
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

neighbor remote-as

`neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} remote-as <AS-NUMBER>`
`no neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} remote-as <AS-NUMBER>`

Description

Creates a peer, initiates the connection to the peer, and adds an entry to the BGP neighbor table. Specifies a neighbor with an autonomous system (AS) number that identifies the neighbor as internal to the local autonomous system. Otherwise, the neighbor is considered as external. By default, neighbors that are defined using this command, exchange only unicast address prefixes.

The **no** form of this command disables the peer session and deletes the peer information.

Parameter	Description
<IP-ADDR>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.

Parameter	Description
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.
`remote-as` *<AS-NUMBER>*	Specifies a 4-byte AS number in asplain format (z), or asdot format (x.y), where z is a number from 1 to 4294967295 and x and y are 16-bit numbers in the range 0 to 65535.

Usage

The configured peer AS number is compared with the AS number received in the open message and a peer session is initiated only if both the AS numbers match.

Examples

```
switch(config-bgp)# neighbor 1.1.1.1 remote-as 1
switch(config-bgp)# no neighbor 1.1.1.1 remote-as 1
```

```
switch(config-bgp)# neighbor pg remote-as 1
switch(config-bgp)# no neighbor pg remote-as 1
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

neighbor remove-private-AS

`neighbor {`*<IP-ADDR>*`|`*<IFNAME>*`|`*<PEER-GROUP-NAME>*`} remove-private-AS`
`no neighbor {`*<IP-ADDR>*`|`*<IFNAME>*`|`*<PEER-GROUP-NAME>*`} remove-private-AS`

Description

This command forces the BGP neighbor to drop the private AS numbers. When the outbound update contains a sequence of private AS numbers, this sequence is dropped. If the command is configured for peer-group, then all the neighbors that are part of peer-group will remove the private-AS before sending the BGP update message.

The **no** form of this command allows the private-AS number to be carried in BGP update message.

Parameter	Description
<IP-ADDRESS>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.

Examples

Forcing the BGP neighbor to drop the private AS numbers:

```
switch(config-bgp)# neighbor 1.1.1.1 remove-private-AS
```

Allowing the private AS number to be carried in the BGP update message:

```
switch(config-bgp)# no neighbor 1.1.1.1 remove-private-AS
```

Forcing the BGP neighbor to drop the private AS numbers:

```
switch(config-bgp)# neighbor PG remove-private-AS
```

Allowing the private AS number to be carried in the BGP update message:

```
switch(config-bgp)# no neighbor PG remove-private-AS
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400	`config-bgp`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300 9300S 10000		

neighbor remote-as

```
neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} remote-as <AS-NUMBER>
no neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} remote-as <AS-NUMBER>
```

Description

Creates a peer, initiates the connection to the peer, and adds an entry to the BGP neighbor table. Specifies a neighbor with an autonomous system (AS) number that identifies the neighbor as internal to the local autonomous system. Otherwise, the neighbor is considered as external. By default, neighbors that are defined using this command, exchange only unicast address prefixes.

The **no** form of this command disables the peer session and deletes the peer information.

Parameter	Description
<IP-ADDR>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.
`remote-as` *<AS-NUMBER>*	Specifies a 4-byte AS number in asplain format (z), or asdot format (x.y), where z is a number from 1 to 4294967295 and x and y are 16-bit numbers in the range 0 to 65535.

Usage

The configured peer AS number is compared with the AS number received in the open message and a peer session is initiated only if both the AS numbers match.

Examples

```
switch(config-bgp)# neighbor 1.1.1.1 remote-as 1
switch(config-bgp)# no neighbor 1.1.1.1 remote-as 1
```

```
switch(config-bgp)# neighbor pg remote-as 1
switch(config-bgp)# no neighbor pg remote-as 1
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

neighbor replace-rt-asn

```
neighbor {<IP-ADDR>|<IPv6-ADDR>|<PEER-GROUP-NAME>} replace-rt-asn
no neighbor {<IP-ADDR>|<IPv6-ADDR>|<PEER-GROUP-NAME>} replace-rt-asn
```

Description

This command enables the BGP router to replace the ASN portion of the route target in the received BGP update message with the local AS number. It facilitates usage of auto RT on all BGP peers across EVPN fabrics.

This command can be used if VNIs are same across fabric and **route-target auto** is being used. When **replace-rt-asn** and a route-map with **set extcommunity rt** applied in the "IN" direction on a neighbor, **replace-rt-asn** will override the 2-byte of extended-community with local-AS.
The bgp connection must be reset using **clear bgp** to allow this configuration to take effect.

Parameter	Description
<IP-ADDR>	Specifies the neighbor's IP address in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.
<IPv6-ADDR>	Specifies the neighbor's IP address in IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported for BGP unnumbered.
`no`	Restores default behavior.

Examples

Specifying the neighbor's IP address:

```
switch(config-bgp-l2vpn-evpn)# neighbor 1.1.1.1 replace-rt-asn
```

Specifying a peer group name:

```
switch(config-bgp-l2vpn-evpn)# neighbor PG replace-rt-asn
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp-l2vpn-evpn`	Administrators or local user group members with execution rights for this command.

neighbor route-map

`neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} route-map <MAP-NAME> {in|out}`
`no neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} route-map <MAP-NAME> {in|out}`

Description

This command applies a route map to incoming or outgoing routes. It configures the route map for modifying the default attributes of the route.

When both peer group and neighbor configuration have route maps associated, then the following configuration applies:

- For outbound route maps, peer group configuration will override the configuration of the neighbor.
- For inbound route maps, neighbor configuration will override the peer group configuration.

The **no** form of this command removes a route map.

Parameter	Description
<IP-ADDRESS>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.
<MAP-NAME>	Specifies the name of the route map.
`in|out`	Sets the route map policy to apply to either the received routes from the neighbor (`in`) or the advertised routes to the neighbor (`out`).

Examples

```
switch(config-bgp-ipv4-uc)# neighbor 1.1.1.1 route-map HPE in
switch(config-bgp-ipv4-uc)# no neighbor 1.1.1.1 route-map HPE in
```

```
switch(config-bgp-ipv4-uc)# neighbor PG route-map HPE in
switch(config-bgp-ipv4-uc)# no neighbor PG route-map HPE in
```

```
switch(config)# route-map Rmap permit seq 10
switch(config-route-map-Rmap-10)# match metric 100
switch(config-route-map-bgp-10)# router bgp 100
switch(config-bgp-ipv4-uc)# neighbor 1.1.1.1 remote-as 100
switch(config-bgp-ipv4-uc)# neighbor 1.1.1.1 route-map Rmap out
```

Configuring inbound route maps in L2VPN EVPN address family.

```
switch(config)# router bgp 100
switch(config-bgp)# neighbor 2.1.1.1 remote-as 100
switch(config-bgp)# neighbor 2.1.1.1 update-source loopback 1
switch(config-bgp)# address-family l2vpn evpn
switch(config-bgp-l2vpn-evpn)# neighbor 2.1.1.1 activate
switch(config-bgp-l2vpn-evpn)# neighbor 2.1.1.1 route-map Rmap in
switch(config-bgp-l2vpn-evpn)# neighbor 2.1.1.1 send-community extended
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority

6300, except for S3L75A, S3L76A, S3L77A
6400
8100
8325
8325H
8325P
8360
9300
9300S
10000
	`config-bgp-ipv4-uc`
	`config-bgp-ipv6-uc`
	`config-bgp-l2vpn-evpn`
		Administrators or local user group members with execution rights for this command.

neighbor route-reflector-client

```
neighbor  {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} route-reflector-client
no neighbor  {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} route-reflector-client
```

Description

This command configures the router as a BGP route reflector and the specified peer as its client. The **no** form of this command disables this function.

Parameter	Description
<IP-ADDRESS>	Specifies the neighbor IP address in the IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255, or in the IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F. **NOTE:** Pv6 MP-BGP peering must not be used for L2VPN EVPN address family, because VXLAN tunnel interface does not support IPv6 addresses.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.

Examples

```
switch(config-bgp-ipv4-uc)# neighbor 1.1.1.1 route-reflector-client
switch(config-bgp-ipv4-uc)# no neighbor 1.1.1.1 route-reflector-client
```

```
switch(config-bgp-ipv4-uc)# neighbor PG route-reflector-client
switch(config-bgp-ipv4-uc)# no neighbor PG route-reflector-client
```

```
switch(config-bgp-l2vpn-evpn)# neighbor 1.1.1.1 route-reflector-client
switch(config-bgp-l2vpn-evpn)# no neighbor 1.1.1.1 route-reflector-client
```

```
switch(config-bgp-l2vpn-evpn)# neighbor PG route-reflector-client
switch(config-bgp-l2vpn-evpn)# no neighbor PG route-reflector-client
```

```
switch(config-bgp-l2vpn-evpn)# neighbor 2001:0db8:85a3::8a2e:0370:7334 route-
reflector-client
switch(config-bgp-l2vpn-evpn)# no neighbor 2001:0db8:85a3::8a2e:0370:7334 route-
reflector-client
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority

6300, except for S3L75A, S3L76A, S3L77A
6400
8100
8325
8325H
8325P
8360
9300
9300S
10000
	`config-bgp-ipv4-uc`
	`config-bgp-ipv6-uc`
	`config-bgp-l2vpn-evpn`
		Administrators or local user group members with execution rights for this command.

neighbor send-community

`neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} send-community [standard | extended]`
`no neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} send-community [standard | extended]`

Description

This command allows community values to be sent to a specific neighbor. When this command is configured for the peer-group, then all the neighbors that are part of peer-group will send the community values to the peers. The parameters `standard` and `extended` send only the respective community numbers. When the command is issued without either of these parameters, both standard and extended communities will be sent to the neighbor.

The **no** form of this command will not allow the neighbor to send community values to the specific neighbors that are part of peer-group.

Parameter	Description
<IP-ADDR>	Specifies the neighbor IP address in the IPv4 format (**x.x.x.x**),

Parameter	Description
	where **x** is a decimal number from 0 to 255, or in the IPv6 format (**xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx**), where **x** is a hexadecimal number from 0 to F. **NOTE:** IPv6 MP-BGP peering must not be used for L2VPN EVPN address family, because VXLAN tunnel interface does not support IPv6 addresses.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.

Examples

```
switch(config-bgp-ipv4-uc)# neighbor 1.1.1.1 send-community standard
switch(config-bgp-ipv4-uc)# no neighbor 1.1.1.1 send-community standard
```

```
switch(config-bgp-ipv4-uc)# neighbor 1.1.1.1 send-community extended
switch(config-bgp-ipv4-uc)# no neighbor 1.1.1.1 send-community
```

```
switch(config-bgp-ipv4-uc)# neighbor PG send-community standard
switch(config-bgp-ipv4-uc)# no neighbor PG send-community standard
```

```
switch(config-bgp-ipv4-uc)# neighbor PG send-community extended
switch(config-bgp-ipv4-uc)# no neighbor PG send-community
```

```
switch(config-bgp-l2vpn-evpn)# neighbor 1.1.1.1 send-community standard
switch(config-bgp-l2vpn-evpn)# no neighbor 1.1.1.1 send-community standard
```

```
switch(config-bgp-l2vpn-evpn)# neighbor 1.1.1.1 send-community extended
switch(config-bgp-l2vpn-evpn)# no neighbor 1.1.1.1 send-community
```

```
switch(config-bgp-l2vpn-evpn)# neighbor PG send-community standard
switch(config-bgp-l2vpn-evpn)# no neighbor PG send-community standard
```

```
switch(config-bgp-l2vpn-evpn)# neighbor PG send-community extended
switch(config-bgp-l2vpn-evpn)# no neighbor PG send-community
```

```
switch(config-bgp-l2vpn-evpn)# neighbor 2001:0db8:85a3::8a2e:0370:7334 send-
community extended
switch(config-bgp-l2vpn-evpn)# no neighbor 2001:0db8:85a3::8a2e:0370:7334 send-
community
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority

6300, except for S3L75A, S3L76A, S3L77A
6400
8100
8325
8325H
8325P
8360
9300
9300S
10000
	`config-bgp-ipv4-uc`
	`config-bgp-ipv6-uc`
	`config-bgp-l2vpn-evpn`
		Administrators or local user group members with execution rights for this command.

neighbor shutdown

```
neighbor   {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} shutdown
no neighbor   {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} shutdown
```

Description

This command disables the peer session, terminates any active session for the specified neighbor or peer group, and removes all associated routing information. This action can cause the sudden termination of many peering sessions.

The **no** form of this command enables the peer session for the specified neighbor.

Parameter	Description
<IP-ADDRESS>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.

Usage

Sessions are gracefully shut down when graceful-shutdown is enabled. Enter the `neighbor graceful-shutdown` command to enable graceful-shutdown. If graceful-shutdown is configured without delay or local-preference, the default values are used.

Examples

Disable the peer session, terminate any active session for the specified neighbor or peer group, and removing all associated routing information. Then enable the peer session for the specified neighbor:

```
switch(config-bgp)# neighbor 1.1.1.1 shutdown
switch(config-bgp)# no neighbor 1.1.1.1 shutdown
```

Disable the peer session, terminate any active session for the specified neighbor or peer group, and removing all associated routing information. Then enable the peer session for the specified neighbor:

```
switch(config-bgp)# neighbor pg shutdown
switch(config-bgp)# no neighbor pg shutdown
```

Establish a peer graceful shutdown:

```
switch(config-bgp)# neighbor 1.1.1.1 graceful-shutdown 10
switch(config-bgp)# neighbor 1.1.1.1 shutdown
Peer 1.1.1.1 will be shutdown in 10 seconds. Do not modify BGP configurations
until the shutdown process completes.
```

Enter an unsuccessful peer session because the gshut timer is running:

```
switch(config-bgp)# no neighbor 1.1.1.1 shutdown
Peer 1.1.1.1 is being gracefully shutdown. Peer cannot be brought up until the
shutdown process completes.
```

Enter an successful peer session because the gshut timer is running:

```
switch(config-bgp)# no neighbor pg shutdown
```

Reset the timer to the default value:

```
switch(config-bgp)# no neighbor pg graceful-shutdown 10
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.16	Added graceful shutdown examples and configuration.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100	config-bgp	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325 8325H 8325P 8360 9300 9300S 10000		

neighbor slow-peer

```
neighbor {<IP-ADDR>|<IPV6-ADDR>|<IFNAME>} slow-peer
no neighbor {<IPV4-ADDR>|<IPV6-ADDR>|<IFNAME>} slow-peer
```

Description

Slow peer configuration moves the peer from its normal update group to a slow update group. Removing slow peers allows the normal update group to function without being slowed down and enables faster convergence. Slow peer configuration is applicable only to explicitly configured BGP neighbors. All configured slow peers will be grouped into a slow update group. This group will converge based on the slowest group peer. Dynamic detection and splitting isn't supported. Peers can be statically configured as slow.

The **no** form of this command removes the slow peer configuration.

Parameter	Description
<IP-ADDR>	Specifies a neighbor IP address.
<IPV6-ADDR>	Specifies a neighbor IPv6 address.

Examples

Configuring a neighbor with an IP address of **1.1.1.1** as a slow peer:

```
switch(config-bgp)# neighbor 1.1.1.1 slow-peer
```

Removing the slow peer configuration for a neighbor with an IP address of **1.1.1.1**:

```
switch(config-bgp)# no neighbor 1.1.1.1 slow-peer
```

Configuring a neighbor with an IPv6 address of **2001::2** as a slow peer:

```
switch(config-bgp)# neighbor 2001::2 slow-peer
```

Removing the slow peer configuration for a neighbor with an IPv6 address of **2001::2**:

```
switch(config-bgp)# no neighbor 2001::2 slow-peer
```

Command History

Release	Modification
10.13	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

neighbor soft-reconfiguration inbound

```
neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} soft-reconfiguration inbound
no neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} soft-reconfiguration inbound
```

Description

This command enables you to generate inbound updates from a neighbor and change and activate BGP policies without clearing the BGP session. Changes in BGP policies require the BGP session to be cleared which can have a large negative impact on network operations.

The **no** form of this command disables this setting.

Parameter	Description
<IP-ADDRESS>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.

Usage

- To perform inbound soft reconfiguration, the BGP speaker must store all received route updates, regardless of the current inbound policy.
- When inbound soft reconfiguration is enabled, the stored updates are processed by the new policy configuration to create new inbound updates.

Examples

```
switch(config-bgp-ipv4-uc)# neighbor 1.1.1.1 soft-reconfiguration inbound
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp-ipv4-uc` `config-bgp-ipv6-uc`	Administrators or local user group members with execution rights for this command.

neighbor timers

```
neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} timers <KEEPALIVE> <HOLDTIME>
no neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} timers <KEEPALIVE> <HOLDTIME>
```

Description

This command sets the timers for a specific BGP neighbor or peer group. When the timer is applied to peer-group then all the neighbors that are part of peer-group will inherit the value configured.

The **no** form of this command clears the timers for a specific BGP neighbor or peer group.

Parameter	Description
<IP-ADDRESS>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.
<KEEPALIVE>	Specifies the Keep-Alive timer value for the neighbor. Default: 60 seconds. Range: 0-65535.
<HOLDTIME>	Specifies the Hold-timer value. Default: 180 seconds. Range: 0-65535.

Examples

```
switch(config-bgp)# neighbor 1.1.1.1 timers 120 360
switch(config-bgp)# no neighbor 1.1.1.1 timers 120 360
```

```
switch(config-bgp)# neighbor pg timers 120 360
switch(config-bgp)# no neighbor pg timers 120 360
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

neighbor timers connect-retry

```
neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} timers connect-retry <RETRYTIME>
no neighbor {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} timers connect-retry <RETRYTIME>
```

Description

This command sets the connect-retry timer for a specific BGP neighbor or peer group.

The **no** form of this command clears the connect-retry timers for a specific BGP neighbor or peer group.

Parameter	Description
<IP-ADDR>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.
<RETRYTIME>	Specifies the retry timer value.

Usage

When the timer is applied to peer-group, then all the neighbors that are part of peer-group inherit the value configured. A connect-retry timer configured for a specific peer takes precedence over that configured for the peer group of this peer. In addition, a connect-retry timer configured for a specific peer or peer group takes precedence over that configured on the global level.

Examples

Setting the connect-retry timer to 50:

```
switch(config-bgp)# neighbor 1.1.1.1 timers connect-retry 50
```

Removing the connect-retry timer:

```
switch(config-bgp)# no neighbor 1.1.1.1 timers connect-retry 50
```

Setting the connect-retry timer to 80:

```
switch(config-bgp)# neighbor pg timers connect-retry 80
```

Removing the connect-retry timer:

```
switch(config-bgp)# no neighbor pg timers connect-retry 80
```

Command History

Release	Modification
10.16	Command introduced.

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-bgp	Administrators or local user group members with execution rights for this command.

neighbor ttl-security-hops

```
neighbor  {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} ttl-security-hops <HOP-COUNT>
no neighbor  {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} ttl-security-hops <HOP-COUNT>
```

Description

This command enables BGP to establish connection with external peers residing on networks that are not directly connected. By enabling this feature, the received TTL from a BGP peer is compared with the difference "255 - *hop-count*". BGP messages coming with a TTL less than this value are not accepted. BGP peering will not be established if the TTL in the session establishment is received with a lower value. Also, by enabling this feature the router will send BGP packets with TTL value of 255 to the neighbor. For a neighbor, either TTL security or `ebgp-multihop` can be configured, not both together. If there are multiple paths to reach the node, then the hop count should be configured considering the longest route.

The **no** form of this command disables the peer ttl-security-hop feature.

Parameter	Description
<IP-ADDR>	Specifies an IP address.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.
<HOP-COUNT>	Specifies the hop count to reach the neighbor for the eBGP session. Range: 1-255.

Examples

```
switch(config-bgp)# neighbor 1.1.1.1 ttl-security-hops 10
switch(config-bgp)# no neighbor 1.1.1.1 ttl-security-hops
```

```
switch(config-bgp)# neighbor pg ttl-security-hops 5
switch(config-bgp)# no neighbor pg ttl-security-hops
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

neighbor update-source

```
neighbor {<IP-ADDRESS>|<PEER-GROUP-NAME>}
     update-source {<IPv4>|<IPv6> | loopback <NUMBER>}
no neighbor {<IP-ADDRESS>|<PEER-GROUP-NAME>}
     update-source [<IPv4>|<IPv6> | loopback <NUMBER>]
```

Description

This command specifies the source address to reach the neighbor.

An iBGP connection can occur as long as there is a TCP/IP path between the routers. If multiple paths exist between the iBGP routers, using a loopback interface as the neighbor address can add stability to the network. With this command, stability can be achieved by providing the loopback interface address as the source address of the TCP/IP session.

The **no** form of this command negates the route updates of the neighbor.

Parameter	Description
<IP-ADDRESS>	Specifies an IP address.
<PEER-GROUP-NAME>	Specifies a peer group.
<IPV4>	Specifies an interface by IPv4 address.
<IPV6>	Specifies an interface by IPv6 address.
`loopback` *<NUMBER>*	Specifies a loopback interface number.

Examples

```
switch(config-bgp)# neighbor 1.1.1.1 update-source loopback 1
switch(config-bgp)# no neighbor 1.1.1.1 update-source
```

```
switch(config-bgp)# neighbor PG update-source loopback 1
switch(config-bgp)# no neighbor PG update-source
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325	`config-bgp`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325H 8325P 8360 9300 9300S 10000		

neighbor weight

```
neighbor  {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} weight <WEIGHT-VALUE>
no neighbor  {<IP-ADDR>|<IFNAME>|<PEER-GROUP-NAME>} weight <WEIGHT-VALUE>
```

Description

This command assigns a weight to a neighbor connection. When the weight is applied to a peer-group then all the neighbors that are part of the peer-group will inherit the value configured.

The **no** form of this command removes a weight assignment.

Parameter	Description
<IP-ADDR>	Specifies an IP address.
<IFNAME>	Specifies the unnumbered interface name of the neighbor.
<PEER-GROUP-NAME>	Specifies a peer group. This is not supported on BGP unnumbered.
<WEIGHT-VALUE>	Specifies the weigh to be associated with the routes received from the neighbor. Range: 0-65535.

Examples

```
switch(config-bgp)# neighbor 1.1.1.1 weight 500
switch(config-bgp)# no neighbor 1.1.1.1 weight
```

```
switch(config-bgp)# neighbor pg weight 600
switch(config-bgp)# no neighbor pg weight
```

Command History

Release	Modification
10.16	Added support for BGP unnumbered.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

network

```
network {[<IPV4/MASK> | <IPV6/MASK>]} [route-map <ROUTE-MAP-NAME>]
no network {[<IPV4/MASK> | <IPV6/MASK>]} [route-map <ROUTE-MAP-NAME>]
```

Description

This command specifies the networks to be advertised by the Border Gateway Protocol (BGP) routing processes.

The **no** form of this command removes an entry from the routing table.

Parameter	Description
<IPV4/MASK>	Specifies the IPv4 network with mask. For example: 1.1.1.1/24
<IPV6/MASK>	Specifies the IPv6 network with mask. For example: 2001:0db8:85a3::8a2e:0370:7334/24
`route-map` *<ROUTE-MAP-NAME>*	Optional parameter. Specifies a route map to apply to the prefixes advertised by this specific network statement.

Usage

- This command is used to advertise prefixes currently installed in the routing table into the BGP table.
- Use the `route-map` keyword to apply the specified route map to network advertisements. The mask length as configured in the network statement must match the mask length of prefixes in the routing table.

Examples

```
switch(config-bgp-ipv4-uc)# network 11.11.11.0/24
switch(config-bgp-ipv4-uc)# no network 11.11.11.0/24
```

```
switch(config-bgp-ipv6-uc)# network 2001:0db8:85a3::8a2e:0370:7334/24
switch(config-bgp-ipv6-uc)# no network 2001:0db8:85a3::8a2e:0370:7334/24
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp-ipv4-uc` `config-bgp-ipv6-uc`	Administrators or local user group members with execution rights for this command.

redistribute

```
redistribute {connected|ospf <process ID> |host-routes|static|local loopback} [route-map <ROUTE-MAP-NAME>]
no redistribute {connected|ospf <process ID> |host-routes|static|local loopback} [route-map <ROUTE-MAP-NAME>]
```

Description

This command specifies routes to import into BGP. This command causes routes from the specified protocol to be considered for redistribution into BGP.

The **no** form of this command specifies no redistribution into BGP.

Parameter	Description
`connected`	Redistributes directly attached networks (directly attached to the subnet or host).
`host-routes`	Redistributes ARP entries. ARP entries learned through EVPN or over VXLAN tunnels are not redistributed.
`ospf`	Redistributes Open Shortest Path First (OSPFv2) routes. It is optional to mention the process ID. Range: <1-65535>
`static`	Redistributes statically configured routes .
`local loopback`	Performs the following functions: ■ Redistributes local routes on loopback interfaces. ■ For EVPN enabled VRFs, it advertises the IP address of loopback interfaces as a EVPN Type-5 prefix route.
`route-map <ROUTE-MAP-NAME>`	Optional. Specifies a route map to match for redistribution.

Usage

- If a route map is specified, then routes that pass the match clause specified in the route map will be imported into the BGP peer Routing Information Base (RIB).
- Route-maps must be configured prior to being referenced in redistribution statements.
- Redistribute connected is required to redistribute connected subnet even if redistribute local loopback is already configured.

Examples

Redistribute directly attached networks:

```
switch(config-bgp-ipv4-uc)# redistribute connected
switch(config-bgp-ipv4-uc)# no redistribute connected
```

Redistributing local routes on loopback interfaces:

```
switch(config-bgp-ipv4-uc)# redistribute local loopback
switch(config-bgp-ipv4-uc)# no redistribute local loopback
```

Redistributing ARP entries:

```
switch(config-bgp-ipv4-uc)# redistribute host-routes route-map map1
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.13	Added **route-map** support for **host-routes**.
10.12	Added support for the `host-routes` parameter.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300	`config-bgp-ipv4-uc` `config-bgp-ipv6-uc`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
9300S 10000		

router bgp

```
router bgp <AS-NUMBER>
no router bgp <AS-NUMBER>
```

Description

This command configures the BGP instance on the router, configures the AS (Autonomous System) the router belongs to, and enters into the BGP router configuration mode. Only a single BGP AS number can be assigned for the entire system.

The **no** form of the command deletes the BGP instance from the router.

Parameter	Description
AS-NUMBER	Specifies a 4-byte AS number in the range 1-4294967295 in integer format or from 0.1-65535.65535 in dotted format.

Examples

Configuring the BGP instance with the AS number:

```
```

Deleting BGP configurations:

```
switch(config)# no router bgp 100
This will delete all BGP configurations on this device.
Continue (y/n)?
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H	config	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325P 8360 9300 9300S 10000		

show bgp

```
show bgp [{vrf <VRF-NAME>|all-vrf}] [{ipv4 unicast|ipv6 unicast
     |all}] [vsx-peer][update-group [<INDEX>]]
show bgp l2vpn evpn
```

Description

This command shows entries in the BGP routing table.

Parameter	Description
`ipv4`	Selects the IPv4 address family.
`ipv6`	Selects the IPv6 address family.
`unicast`	The subaddress family identifier.
`vrf <VRF-NAME>`	Select to display information by VRFs by specifying the VRF name.
`all-vrf`	Select to display the BGP summary information for all VRFs and address-families.
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.
`l2vpn evpn`	Shows the information for L2VPN EVPN address family. This parameter only applies to 8325 Series switches.
`update-group [<INDEX>]`	Select to display information about update-group and peers which are member of that update-group for given afi/safi. Specifying an integer defines the desired update-group index.

Examples

Showing BGP routing table information for VRF 1 IPv4 unicast:

```
switch# show bgp vrf v1 ipv4 unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : v1
Local Router-ID 9.0.0.1

    Network            Nexthop        Metric     LocPrf     Weight Path
-------------------------------------------------------------------------
*>e 9.0.0.0/24         9.0.0.2        0          100        0      65534.65535 3.4
```

```
18.54934 3574.8570 5.6 ?
*>e 100.0.0.0/24       9.0.0.2         0          100         0      200 ?
*>e 100.0.1.0/24       9.0.0.2         0          100         0      200 ?
*>e 100.0.2.0/24       9.0.0.2         0          100         0      200 ?
*>e 100.0.3.0/24       9.0.0.2         0          100         0      200 ?
*ae 100.0.3.0/24       9.0.0.3         0          100         0      200 ?
Total number of entries 6
```

Showing BGP routing table information for L2VPN EVPN:

```
switch# show bgp l2vpn evpn

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete
EVPN Route-Type 2 prefix: [2]:[ESI]:[EthTag]:[MAC]:[OrigIP]
EVPN Route-Type 3 prefix: [3]:[EthTag]:[OrigIP]

VRF : default
Router-ID not configured

     Network                                                    Nexthop
Metric LocPrf Weight Path
-------------------------------------------------------------------------------
----------------------
Route Distinguisher: 10.1.1.54:32967    (L2VNI 30000)
*>   [2]:[0]:[0]:[00:06:f6:3f:e3:c1]:[]                        1.1.1.20        0
   100    32768  i
*>   [2]:[0]:[0]:[8c:60:4f:f2:f5:41]:[]                        1.1.1.10        0
   100    0      i
*>   [3]:[0]:[1.1.1.1]                                         0.0.0.0         0
   100    0      ?
Total number of entries 3
```

BGP routing information for a network that includes both IPv4 and IPv6 addresses.

```
switch# show bgp l2vpn evpn vtep 1920:1680:1:1::4 vni 1001001
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete
EVPN Route-Type 2 prefix: [2]:[ESI]:[EthTag]:[MAC]:[OrigIP]
EVPN Route-Type 3 prefix: [3]:[EthTag]:[OrigIP]
EVPN Route-Type 5 prefix: [5]:[ESI]:[EthTag]:[IPAddrLen]:[IPAddr]
VRF : default
Local Router-ID 192.168.1.1
Network                                                Nexthop         Metric
LocPrf    Weight   Path
-------------------------------------------------------------------------------
--------------------------
Route Distinguisher: 192.168.1.4:1001      (L2VNI 1001001)
*>i [2]:[0]:[0]:[00:aa:bb:cc:11:01]:[100.1.1.1]            1920:1680:1:1::4
0          100        0       ?
* i [2]:[0]:[0]:[00:aa:bb:cc:11:01]:[100.1.1.1]            1920:1680:1:1::4
0          100        0       ?
*>i [2]:[0]:[0]:[00:aa:bb:cc:11:01]:[1000:1:1:1::1]        1920:1680:1:1::4
0          100        0       ?
* i [2]:[0]:[0]:[00:aa:bb:cc:11:01]:[1000:1:1:1::1]        1920:1680:1:1::4
0          100        0       ?
*>i [2]:[0]:[0]:[00:aa:bb:cc:11:01]:[fe80:0:1::1]          1920:1680:1:1::4
```

```
0          100        0       ?
* i [2]:[0]:[0]:[00:aa:bb:cc:11:01]:[fe80:0:1::1]          1920:1680:1:1::4
0          100        0       ?
*>i [3]:[0]:[1920:1680:1:1::4]                             1920:1680:1:1::4
0          100        0       ?
* i [3]:[0]:[1920:1680:1:1::4]                             1920:1680:1:1::4
0          100        0       ?
Total number of entries 8
```

```
switch# show bgp l2vpn evpn neighbors 1920:1680:1:1::8
Codes: ^ Inherited from peer-group, * Dynamic Neighbor
VRF : default
BGP Neighbor 1920:1680:1:1::8 (Internal)
Description          : RR peer-group^
Peer-group           : RRv6
Remote Router Id     : 192.168.1.8        Local Router Id     : 192.168.1.1
Remote AS            : 65001              Local AS            : 65001
Remote Port          : 42423              Local Port          : 179
State                : Established        Admin Status        : Up
Conn. Established    : 5                  Conn. Dropped       : 4
Passive              : No                 Update-Source       : loopback0^
Cfg. Hold Time       : 180                Cfg. Keep Alive     : 60
Neg. Hold Time       : 180                Neg. Keep Alive     : 60
Up/Down Time         : 06h:46m:13s        Connect-Retry Time  : 120
Local-AS Prepend     : No                 Alt. Local-AS       : 0
BFD                  : Disabled           Slow Peer           : Yes
Password             :
Last Err Sent        : No Error
Last SubErr Sent     : No Error
Last Err Rcvd        : No Error
Last SubErr Rcvd     : No Error
Graceful-Restart     : Enabled            Gr. Restart Time    : 120
Gr. Stalepath Time   : 300                Remove Private-AS   : No
TTL                  : 255                Local Cluster-ID    :
Weight               : 0                  Fall-over           : No
Confederation-Peers : No

Message statistics        Sent     Rcvd
-------------------      -----    -----
Open                        8        7
Notification                3        1
Updates                 20730    91332
Keepalives               1153      952
Route Refresh               0        0
Total                   21894    92292

Capability                          Advertised      Received
---------------------------         -----------     ----------
Route Refresh                       Yes             Yes
Graceful Restart                    Yes             Yes
Add-Path                            No              No
Four Octet ASN                      Yes             Yes
Address family IPv4 Unicast         No              No
Address family IPv6 Unicast         No              No
Address family VPNv4 Unicast        No              No
Address family L2VPN EVPN           Yes             Yes
Address Family : L2VPN EVPN
-----------------------------
Rt. Reflect. Client : No                Send Community    : extended^
Allow-AS in         : 0                 Advt. Interval    : 30
Max. Prefix         : 64000             Soft Reconfig In  :
```

```
Nexthop-Self        :                       Default-Originate :
Cfg. Add-Path       :
Neg. Add-Path       :
Routemap In         :
Routemap Out        :
ORF type            : Prefix-list
ORF capability      :
```

Showing all BGP update groups:

```
switch# show bgp all update-group

VRF : default

BGP Update-Group 1
    Address Family  : ipv4-unicast        Peering Type  : internal
    Peer Count      : 5                   Slow Group    : No

    Members
    -------
    110.162.100.221, 120.221.221.221, 130.121.21.111, 140.131.131.131,
    150.100.101.2

BGP Update-Group 2
    Address Family  : ipv4-unicast        Peering Type  : external
    Peer Count      : 2                   Slow Group    : Yes

    Members
    -------
    20.1.1.1, 30.1.1.1

BGP Update-Group 3
    Address Family  : ipv6-unicast        Peering Type  : external
    Peer Count      : 1                   Slow Group    : No

    Members
    -------
    20ab::cd:08
```

Command History

Release	Modification
10.14	Replaced IPv4 with IP in commands. IPv4 is deprecated.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100	Operator (>) or Manager (#)	

Platforms	Command context	Authority
8325 8325H 8325P 8360 9300 9300S 10000		

show bgp

```
show bgp
   [vrf <VRF-NAME>] {ipv4 unicast <A.B.C.D/M>}|{ipv6 unicast <X::Y/M>}
   [vsx-peer]
   l2vpn evpn
      EVPN Route-Type 1 prefix: [1]:[ESI]:[EthTag]
      RD-[ROUTE_TYPE]:[ESI]:[EthTag]:[MAC]:[OrigIP]
      RD-[ROUTE_TYPE]:[EthTag]:[OrigIP]
      RD-[ROUTE_TYPE]:[ESI]:[EthTag]:[IPAddrLen]:[IPAddr]
```

Description

This command displays entries in the BGP routing table that are part of the specified network. For EVPN Route-type 2 with MAC only prefix as an input, displays all the prefixes containing the specific MAC address (MAC route, MAC/IP route, Host route).

Parameter	Description
`vrf <VRF-NAME>`	Shows the information for a specified VRF.
`ipv4 unicast <A.B.C.D/M>`	Shows the information for an IPv4 unicast family with an IP prefix (network/length such as 35.0.0.0/8) in the BGP routing table to display.
`ipv6 unicast <X::Y/M>`	Shows the information for an IPv6 unicast family an IPv6 prefix in the BGP routing table to display.
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.
`l2vpn evpn`	Shows the information for L2VPN EVPN address family.
`RD-[ROUTE_TYPE]:[ESI]:[EthTag]:[MAC]:[OrigIP]`	EVPN Route-Type 1 prefix. This parameter is supported on HPE Aruba Networking 8325 and 10000 Switch series only.
`RD-[ROUTE_TYPE]:[ESI]:[EthTag]:[MAC]:[OrigIP]`	EVPN Route-Type 2 prefix.

Parameter	Description
RD-[ROUTE_TYPE]:[EthTag]:[OrigIP]	EVPN Route-Type 3 prefix.
RD-[ROUTE_TYPE]:[ESI]:[EthTag]:[IPAddrLen]:[IPAddr]	EVPN Route-Type 5 prefix.

Examples

Showing the entries in the BGP routing table that are part of an IPv4 unicast network

```
switch# show bgp ipv4 unicast 10.0.0.0/16

VRF : default
BGP Local AS 2         BGP Router-ID 1.1.1.2

   Network           : 10.0.0.0/16              Nexthop             : 1.1.1.1
   Peer              : 1.1.1.1                  Origin              : IGP
   Metric            : 0                        Local Pref          : 100
   Weight            : 0                        Calc. Local Pref    : 100
   Best              : Yes                      Valid               : Yes
   Type              : external                 Stale               : No
   Originator ID     : 0.0.0.0                  Path ID             : 0
   Aggregator ID     :
   Aggregator AS     :
   Atomic Aggregate  :
   RFD Flaps         : 0                        RFD Penalty         : 0

   AS-Path           : 1
   Cluster List      :
   Communities       :
50:100,50:101,50:102,50:103,50:104,50:105,50:106,50:107,50:108,50:109,50:110,50:1
   Extd. Communities :
```

Showing the entries in the BGP routing table that are part of L2VPN EVPN

```
switch# show bgp l2vpn evpn vni 30000
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete
EVPN Route-Type 2 prefix: [2]:[ESI]:[EthTag]:[MAC]:[OrigIP]
EVPN Route-Type 3 prefix: [3]:[EthTag]:[OrigIP]

VRF : default
Router-ID not configured

     Network                                                Nexthop
Metric LocPrf Weight Path
-------------------------------------------------------------------------------
----------------------
Route Distinguisher: 10.1.1.54:32967    (L2VNI 30000)
*>   [2]:[0]:[0]:[00:06:f6:3f:e3:c1]:[]                    1.1.1.20        0
   100    32768  i
*>   [2]:[0]:[0]:[8c:60:4f:f2:f5:41]:[]                    1.1.1.10        0
   100    0      i
*>   [3]:[0]:[1.1.1.1]                                     0.0.0.0         0
   100    0      ?
Total number of entries 3
```

Showing the entries in the BGP routing table for EVPN route-type 2

```
switch# show bgp l2vpn evpn 2:2-[2]:[0]:[0]:[00:50:56:96:6d:6f]:[20.20.1.10]
VRF : default
BGP Local AS 1           BGP Router-id 3.3.3.3
Network           : 2:2-[2]:[0]:[0]:[00:50:56:96:6d:6f]:[20.20.1.10]
Nexthop           : 1.1.1.1
vni               : 2                         vni_type            : L2VNI
Peer              : 2.2.2.2                   Origin              : incomplete
Metric            : 0                         Local Pref          : 100
Weight            : 0                         Calc. Local Pref    : 100
Best              : Yes                       Valid               : Yes
Type              : internal                  Stale               : No
Originator ID     : 1.1.1.1                   Aggregator ID       :
Aggregator AS     :
Atomic Aggregate  :
AS-Path           :
Cluster List      :
Communities       :
Ext-Communities   : RT: 2:2 RT: 10:10 Router MAC: 00:00:00:00:00:11
Network           : 2:2-[2]:[0]:[0]:[00:50:56:96:6d:6f]:[20.20.1.10]
Nexthop           : 1.1.1.1
vni               : 10000                     vni_type            : L3VNI
Peer              : 2.2.2.2                   Origin              : incomplete
Metric            : 0                         Local Pref          : 100
Weight            : 0                         Calc. Local Pref    : 100
Best              : Yes                       Valid               : Yes
Type              : internal                  Stale               : No
Originator ID     : 1.1.1.1                   Aggregator ID       :
Aggregator AS     :
Atomic Aggregate  :
AS-Path           :
Cluster List      :
Communities       :
Ext-Communities   : RT: 2:2 RT: 10:10 Router MAC: 00:00:00:00:00:11
```

Command History

Release	Modification
10.14.1000	8325 and 10000 Switch series can display information for an EVPN Route-Type 1 prefix
10.14	Replaced IPv4 with IP in commands. IPv4 is deprecated.
10.08	Added l2vpn evpn route types
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P	Operator (>) or Manager (#)	

Platforms	Command context	Authority
8360 9300 9300S 10000		

show bgp community

```
show bgp [{ipv4 | ipv6 | ipv4 {vrf <VRF-NAME>}} unicast] community [<VALUE> | <TYPE>]
[vsx-peer]
```

Description

This command shows routes that belong to BGP communities. Optionally you can specify displaying information by a specific community or by VRF.

Parameter	Description
`ipv4`	Shows the information for an IPv4 address family.
`ipv6`	Shows the information for an IPv6 address family.
`unicast`	Shows the information for a subaddress family identifier.
`ipv4 vrf <VRF-NAME>`	Shows the information for a specified VRF.
<VALUE>	Shows the information for a community number. Specify the information in `aa:nn` format.

<TYPE>	Shows a specified community type. Select the following well-known communities, as well as others:
	`internet`
	Advertise the prefix to all BGP neighbors.
	`local-as`
	Do not advertise the prefix outside the sub-AS.
	`no-advertise`
	Do not advertise the prefix to any BGP neighbors.
	`no-export`
	Do not advertise the prefix to any eBGP neighbors.

`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Examples

Showing information for routes that belong to all BGP communities:

```
switch# show bgp ipv4 unicast community
Status codes: s suppressed, d damped, h history, * valid, &gt; best, = multipath,i
internal, e external S Stale, R Removed, a additional-pathsVRF : defaultLocal
Router-ID 9.0.0.1
Network                Nexthop         Metric     LocPrf     Weight Path
-------------------------------------------------------------------------
*&gt;e 9.0.0.0/24      9.0.0.2         0          100        0      200 ?
```

```
*&gt;e 100.0.0.0/24    9.0.0.2        0          100        0      200 ?
*&gt;e 100.0.1.0/24    9.0.0.2        0          100        0      200 ?
*&gt;e 100.0.2.0/24    9.0.0.2        0          100        0      200 ?
*&gt;e 100.0.3.0/24    9.0.0.2        0          100        0      200 ?
*ae 100.0.3.0/24       9.0.0.3        0          100        0      200 ?
                       Total number of entries 6
```

Showing information for routes that belong to the 200:20 BGP community number:

```
switch# show bgp ipv4 unicast community 200:20
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R RemovedOrigin codes: i - IGP, e - EGP, ? -
incomplete
VRF : defaultLocal Router-ID 9.0.0.1

Network              Nexthop        Metric     LocPrf     Weight Path
-------------------------------------------------------------------------
*>e 9.0.0.0/24       9.0.0.2        0          100        0      200 ?
*>e 100.0.0.0/24     9.0.0.2        0          100        0      200 ?
*>e 100.0.1.0/24     9.0.0.2        0          100        0      200 ?
*>e 100.0.2.0/24     9.0.0.2        0          100        0      200 ?
*>e 100.0.3.0/24     9.0.0.2        0          100        0      200 ?
*ae 100.0.3.0/24     9.0.0.3        0          100        0      200 ?
                        Total number of entries 6
```

Showing information for routes that belong to the Internet BGP community type:

```
switch# show bgp ipv4 unicast community internet
Status codes: s suppressed, d damped, h history, * valid, &gt; best, = multipath,
i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete
VRF : default
Local Router-ID 9.0.0.1
Network                Nexthop        Metric     LocPrf     Weight Path
-------------------------------------------------------------------------
*&gt;e 9.0.0.0/24       9.0.0.2        0          100        0      200 ?
*&gt;e 100.0.0.0/24     9.0.0.2        0          100        0      200 ?
*&gt;e 100.0.1.0/24     9.0.0.2        0          100        0      200 ?
*&gt;e 100.0.2.0/24     9.0.0.2        0          100        0      200 ?
*&gt;e 100.0.3.0/24     9.0.0.2        0          100        0      200 ?
*ae 100.0.3.0/24       9.0.0.3        0          100        0      200 ?
Total number of entries 6
```

Showing information for routes that belong to the local-as BGP community type:

```
switch# show bgp ipv4 unicast community local-as
Status codes: s suppressed, d damped, h history, * valid, &gt; best, = multipath,
i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete
VRF : default


Local Router-ID 9.0.0.1
Network                Nexthop        Metric     LocPrf     Weight Path
-------------------------------------------------------------------------
*&gt;e 9.0.0.0/24       9.0.0.2        0          100        0      200 ?
*&gt;e 100.0.0.0/24     9.0.0.2        0          100        0      200 ?
```

```
*&gt;e 100.0.1.0/24    9.0.0.2         0          100         0      200 ?
*&gt;e 100.0.2.0/24    9.0.0.2         0          100         0      200 ?
*&gt;e 100.0.3.0/24    9.0.0.2         0          100         0      200 ?
*ae 100.0.3.0/24       9.0.0.3         0          100         0      200 ?
                       Total number of entries 6
```

Showing information for routes that belong to the no-advertise BGP community type:

```
switch# show bgp ipv4 unicast community no-advertise
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R Removed, a additional-pathsOrigin codes: i -
IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 9.0.0.1
Network            Nexthop         Metric     LocPrf     Weight Path
--------------------------------------------------------------------------
*>e 9.0.0.0/24         9.0.0.2         0          100         0      200 ?
*>e 100.0.0.0/24       9.0.0.2         0          100         0      200 ?
*>e 100.0.1.0/24       9.0.0.2         0          100         0      200 ?
*>e 100.0.2.0/24       9.0.0.2         0          100         0      200 ?
*>e 100.0.3.0/24       9.0.0.2         0          100         0      200 ?
*ae 100.0.3.0/24       9.0.0.2         0          100         0      200 ?
                       Total number of entries 6
```

Showing information for routes that belong to the no-export BGP community type:

```
switch# show bgp ipv4 unicast community no-export
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,i
internal, e external S Stale, R Removed, a additional-pathsOrigin codes: i - IGP,
e - EGP, ? - incomplete

VRF : default
Local Router-ID 9.0.0.1

Network            Nexthop         Metric     LocPrf     Weight Path
--------------------------------------------------------------------------
*>e 9.0.0.0/24         9.0.0.2         0          100         0      200 ?
*>e 100.0.0.0/24       9.0.0.2         0          100         0      200 ?
*>e 100.0.1.0/24       9.0.0.2         0          100         0      200 ?
*>e 100.0.2.0/24       9.0.0.2         0          100         0      200 ?
*>e 100.0.3.0/24       9.0.0.2         0          100         0      200 ?
*ae 100.0.3.0/24       9.0.0.3         0          100         0      200 ?
                       Total number of entries 6
```

Showing information for routes that belong to the gshut BGP community type:

```
switch# show bgp ipv4 unicast community gshut
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
i internal, e external S Stale, R RemovedOrigin codes: i - IGP, e - EGP, ? -
incomplete

VRF : default
Local Router-ID 1.1.1.2

Network            Nexthop          Metric     LocPrf     Weight Path
--------------------------------------------------------------------------
```

```
*>e 1.1.1.0/24          10.1.1.2         0          0          0       2 i
Total number of entries 1

switch#
switch# show bgp ipv6 unicast community gshut

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,i
internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 1.1.1.2

Network            Nexthop          Metric     LocPrf     Weight Path
-------------------------------------------------------------------------
*>e 1::/64             10::2
fe80::98f2:b300:1368:e882
                                0          0          0       2 i
                       Total number of entries 1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show bgp flap-statistics

`show bgp {vrf <VRF_NAME> | all-vrf} [{ipv4 unicast | ipv6 unicast | all}] flap-statistics`

Description

Displays all the flapped and suppressed routes.

Usage

Status of the route with dampening enabled:

- If the route is available, the history flag is unset.
- If route has been flapping, is not suppressed and is withdrawn; the state of the route is `h`

- If route is currently available but is suppressed due to dampening, the state of the route is `d`
- If the route is unsuppressed and currently withdrawn, the state of the route is `h`

Examples

Showing all the flapped and suppressed routes:

```
switch# show bgp all
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 1.1.1.1

Address-family : IPv4 Unicast
------------------------------
    Network            Nexthop         Metric     LocPrf     Weight Path
*>i 2.2.2.0/24         2.2.2.2         0          100        0       ?
*>i 11.1.1.0/24        2.2.2.2         0          100        0       ?
*ai 11.1.1.0/24        2.2.2.3         0          100        0       ?
Total number of entries 3

Address-family : IPv6 Unicast
------------------------------
    Network            Nexthop         Metric     LocPrf     Weight Path
Total number of entries 0
```

```
switch# show bgp ipv4 unicast flap-statistics
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
             i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 20.0.0.1

    Network            Nexthop         Flaps         Reuse      Path
*>e 2.2.2.0/24         20.0.0.2        1          00h:00m:00s   300 ?
 de 3.3.3.0/24         20.0.0.2        2          00h:29m:31s   300 ?
Total number of entries 2
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325 8325H 8325P 8360 9300 9300S 10000		

show bgp ip unicast

```
show bgp ip unicast [<A.B.C.D/M> | protected]
```

Description

This command shows configuration for all BGP routes. Protected Routes will have two Nexthops and the backup path is identified with the status code 'b'

Usage

- Protected Routes will have two Nexthops
- The backup path is identified with the status code 'b'

Examples

Showing BGP IP routes:

```
switch# show bgp ip unicast
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths,
              b backup path
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 200.0.0.1

Network                          Nexthop         Metric     LocPrf       Weight
   Path
Route Distinguisher: 1:0
*>  10.10.10.0/24               0.0.0.0         0          100          0
     ?
* i 10.10.10.0/24               10.10.10.1      0          100          0
     ?
*>i 20.20.20.0/24               10.10.10.1      0          100          0
     ?
*=i 20.20.20.0/24               20.20.20.2      0          100          0
     ?
*=i 20.20.20.0/24               40.40.40.2      0          100          0
     ?
*>  30.30.30.0/24               0.0.0.0         0          100          0
     ?
* i 30.30.30.0/24               30.30.30.1      0          100          0
     ?
*=i 40.40.40.0/24               20.20.20.2      0          100          0
     ?
*>i 40.40.40.0/24               30.30.30.1      0          100          0
     ?
*=i 40.40.40.0/24               40.40.40.2      0          100          0
```

```
     ?
 *>  200.0.0.0/24               0.0.0.0         0          100          0
     i
*>i 201.0.0.0/24                20.20.20.2      0          100          0
     i
*bi 201.0.0.0/24                40.40.40.2      0          100          0
     i
Total number of entries 13
```

Showing protected BGP IP routes:

```
switch# show bgp ip unicast protected
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths,
              b backup path
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 200.0.0.1

Network                         Nexthop         Metric     LocPrf       Weight
   Path
Route Distinguisher: 1:0
*>i 201.0.0.0/24                20.20.20.2      0          100          0
   i
*bi 201.0.0.0/24                40.40.40.2      0          100          0
   i
Total number of entries 2
```

Showing BGP routes for a specific IP address:

```
switch# show bgp ip unicast 201.0.0.0/24
VRF : default
BGP Local AS 1           BGP Router-id 200.0.0.1

Network           : 201.0.0.0/24              Nexthop            : 20.20.20.2
Peer              : 10.10.10.1                Origin             : IGP
Metric            : 0                         Local Pref         : 100
Weight            : 0                         Calc. Local Pref   : 100
Best              : Yes                       Valid              : Yes
Type              : internal                  Stale              : No
Originator ID     : 201.0.0.1
Aggregator ID     :
Aggregator AS     :
Atomic Aggregate  :
RFD Flaps         : 0                         RFD Penalty        : 0
AS-Path           :
Cluster List      : 10.10.10.1
Communities       :
Ext-Communities   :
Backup            : No
Network           : 201.0.0.0/24              Nexthop            : 40.40.40.2
Peer              : 30.30.30.1                Origin             : IGP
Metric            : 0                         Local Pref         : 100
Weight            : 0                         Calc. Local Pref   : 100
Best              : No                        Valid              : Yes
```

```
Type              : internal                  Stale               : No
Originator ID     : 201.0.0.1
Aggregator ID     :
Aggregator AS     :
Atomic Aggregate  :
RFD Flaps         : 0                         RFD Penalty         : 0
AS-Path           :
Cluster List      : 30.30.30.1
Communities       :
Ext-Communities   :
Backup            : Yes
```

Command History

Release	Modification
10.15	Introduces the **protected** option.
10.14.1000	Command introduced.

Command Information

Platforms	Command context	Authority
9300 9300S	Operator (>) or Manager (#)	

show bgp neighbor advertised-routes

```
show bgp [vrf <VRF-NAME>] {ipv4 unicast | ipv6 unicast
      | l2vpn evpn | all} neighbors <IP-ADDRESS>
      advertised-routes [vsx-peer]
show bgp l2vpn evpn neighbors <IP-ADDRESS> advertised-routes
```

Description

Shows all routes that have been advertised to the specified neighbor.

Parameter	Description
`vrf <VRF-NAME>`	Shows the information for a specified VRF.
`ipv4 unicast`	Shows the information for an IPv4 unicast address family.
`ipv6 unicast`	Shows the information for an IPv6 unicast address family.
`l2vpn evpn`	Shows the information for L2VPN EVPN address family.
`all`	Shows the information for all address families and subaddress families.
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Parameter	Description
<IP-ADDRESS>	Shows the information for a neighbor IP address.

Examples

Showing routes that have been advertised to the specified IPv4 unicast neighbor:

```
switch# show bgp ipv4 unicast neighbors 9.0.0.1 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 9.0.0.2

    Network            Nexthop         Metric     LocPrf     Weight Path
------------------------------------------------------------------------
*>e 9.0.0.0/24         9.0.0.2         0          0          0      200 65534.65535
3.4 18.54934 3574.8570 5.6 ?
*>e 100.0.0.0/24       9.0.0.2         0          0          0      200 ?
*>e 100.0.1.0/24       9.0.0.2         0          0          0      200 ?
*>e 100.0.2.0/24       9.0.0.2         0          0          0      200 ?
*>e 100.0.3.0/24       9.0.0.2         0          0          0      200 ?
Total number of entries 5
```

Showing routes that have been advertised to the specified L2VPN EVPN neighbor:

```
switch# show bgp l2vpn evpn neighbor 9.0.0.1 advertised-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete
EVPN Route-Type 2 prefix: [2]:[ESI]:[EthTag]:[MAC]:[OrigIP]
EVPN Route-Type 3 prefix: [3]:[EthTag]:[OrigIP]

VRF : default
Local Router-ID 9.0.0.2

     Network                                                Nexthop         Metric
LocPrf Weight Path
------------------------------------------------------------------------------------
----------------------
Route Distinguisher: 10.1.1.54:32967    (L2VNI 30000)
*>   [2]:[0]:[0]:[00:06:f6:3f:e3:c1]:[]                    1.1.1.20        0      100
  32768  i
*>   [2]:[0]:[0]:[8c:60:4f:f2:f5:41]:[]                    1.1.1.10        0      100
  0      i
Total number of entries 2
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show bgp neighbor paths

```
show bgp [vrf <VRF-NAME>] {ipv4 unicast | ipv6 unicast
     | all} neighbors <IP-ADDRESS> paths [vsx-peer]
show bgp l2vpn evpn neighbors <IP-ADDRESS> paths
```

Description

Shows autonomous system paths learned from the specified neighbor.

Parameter	Description
`vrf <VRF-NAME>`	Shows the information for a specified VRF.
`ipv4 unicast`	Shows the information for an IPv4 unicast address family.
`ipv6 unicast`	Shows the information for an IPv6 unicast address family.
`all`	Shows the information for all address families and subaddress families.
`<IP-ADDRESS>`	Shows the information for a neighbor IP address.
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.
`l2vpn evpn`	Shows the information for L2VPN EVPN address family. This parameter applies only to 8325 series switches.

Examples

Showing autonomous system paths learned from the specified IPv4 unicast neighbor:

```
switch# show bgp ipv4 unicast neighbors 192.168.12.2 paths
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
VRF : default
Local Router-ID 9.0.0.1

    Network            Nexthop          Path
-------------------------------------------
*>e 9.0.0.0/24         9.0.0.2          200 65534.65535 3.4 18.54934 3574.8570 5.6
```

```
*>e 100.0.0.0/24       9.0.0.2          200
*>e 100.0.1.0/24       9.0.0.2          200
*>e 100.0.2.0/24       9.0.0.2          200
*>e 100.0.3.0/24       9.0.0.2          200
Total number of entries 5
```

Showing autonomous system paths learned from the specified L2VPN EVPN neighbor:

```
switch# show bgp l2vpn evpn neighbors 192.168.12.1 paths

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete
EVPN Route-Type 2 prefix: [2]:[ESI]:[EthTag]:[MAC]:[OrigIP]
EVPN Route-Type 3 prefix: [3]:[EthTag]:[OrigIP]

VRF : default
Local Router-ID 9.0.0.2

     Network                                                Nexthop         Path
--------------------------------------------------------------------------------
-
Route Distinguisher: 10.1.1.54:32967    (L2VNI 30000)
*>   [2]:[0]:[0]:[00:06:f6:3f:e3:c1]:[]                 1.1.1.20         100
*>   [2]:[0]:[0]:[8c:60:4f:f2:f5:41]:[]                 1.1.1.10         100
                      Total number of entries 2
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show bgp neighbor received orf-prefix-list

```
show bgp [vrf <VRF-NAME>]{ipv4 unicast | ipv6 unicast | all} neighbors <IP-ADDRESS>
received orf-prefix-list
```

Description

Shows all the prefix lists received from the specified neighbor.

Parameter	Description
`vrf <VRF-NAME>`	Shows the information for a specified VRF.
`ipv4 unicast`	Shows the information for an IPv4 unicast address family.
`ipv6 unicast`	Shows the information for an IPv6 unicast address family.
`all`	Shows the information for all address families and subaddress families.
<IP-ADDRESS>	Shows the information for a neighbor IP address.

Examples

Showing received prefix list from the specified neighbor:

```
switch# show bgp ipv4 unicast neighbors A.B.C.D received orf-prefix-list
Address family: IPv4 Unicast

  ip prefix-list 10.0.0.200: 4 entries

     seq 10 permit 28.119.16.0/24

     seq 15 deny 28.119.19.0/24

     seq 20 permit 28.119.17.0/24

  Address family: IPv6 Unicast

  ip prefix-list 10.0.0.200: 4 entries

       seq 30 permit 2000::/64

       seq 35 deny 3000::/64

       seq 40 permit 4000:0/64
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300	Operator (>) or Manager (#)	

Platforms	Command context	Authority
9300S 10000		

show bgp neighbor received-routes

```
show bgp [vrf <VRF-NAME>] {ipv4 unicast | ipv6 unicast
     | all} neighbors <IP-ADDRESS> received-routes [vsx-peer]
show bgp l2vpn evpn neighbors <IP-ADDRESS> received-routes
```

Description

Shows received routes from the specified neighbor.

Parameter	Description
`vrf <VRF-NAME>`	Shows the information for a specified VRF.
`ipv4 unicast`	Shows the information for an IPv4 unicast address family.
`ipv6 unicast`	Shows the information for an IPv6 unicast address family.
`all`	Shows the information for all address families and subaddress families.
`<IP-ADDRESS>`	Shows the information for a neighbor IP address.
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.
`l2vpn evpn`	Shows the information for L2VPN EVPN address family. This parameter only applies to 8325 Series switches.

Examples

Showing received routes from the specified IPv4 unicast neighbor:

```
switch# show bgp ipv4 unicast neighbors 192.168.12.1 received-routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 9.0.0.2

    Network            Nexthop         Metric     LocPrf     Weight Path
--------------------------------------------------------------------------
*>e 9.0.0.0/24         9.0.0.2         0          0          0      200 65534.65535
3.4 18.54934 3574.8570 5.6 ?
*>e 100.0.0.0/24       9.0.0.2         0          0          0      200 ?
*>e 100.0.1.0/24       9.0.0.2         0          0          0      200 ?
*>e 100.0.2.0/24       9.0.0.2         0          0          0      200 ?
*>e 100.0.3.0/24       9.0.0.2         0          0          0      200 ?
*ae 100.0.3.0/24       9.0.0.2         0          0          0      200 ?
Total number of entries 6
```

Showing received routes from the specified L2VPN EVPN neighbor:

```
switch# show bgp l2vpn evpn neighbors 192.168.12.1 received-routes

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete
EVPN Route-Type 2 prefix: [2]:[ESI]:[EthTag]:[MAC]:[OrigIP]
EVPN Route-Type 3 prefix: [3]:[EthTag]:[OrigIP]

VRF : default
Local Router-ID 9.0.0.2

     Network                                                Nexthop         Metric
LocPrf Weight Path
------------------------------------------------------------------------------------
----------------------
Route Distinguisher: 10.1.1.54:32967    (L2VNI 30000)
*>   [2]:[0]:[0]:[00:06:f6:3f:e3:c1]:[]                 1.1.1.20        0       100
  32768  i
*>   [2]:[0]:[0]:[8c:60:4f:f2:f5:41]:[]                 1.1.1.10        0       100
  0      i
Total number of entries 2
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show bgp neighbor routes

```
show bgp [vrf <VRF-NAME>] {ipv4 unicast | ipv6 unicast
     | all} neighbors <IP-ADDRESS> routes [vsx-peer]
show bgp l2vpn evpn neighbors <IP-ADDRESS>
```

Description

This command shows routes that are received and accepted from the specified neighbor.

Parameter	Description
`vrf <VRF-NAME>`	Shows the information for a specified VRF.
`ipv4 unicast`	Shows the information for an IPv4 unicast address family.
`ipv6 unicast`	Shows the information for an IPv6 unicast address family.
`all`	Shows the information for all address families and subaddress families.
`<IP-ADDRESS>`	Shows the information for a neighbor IP address.
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.
`l2vpn evpn`	Shows the information for L2VPN EVPN address family. This parameter only applies to 8325 Series switches.

Examples

Showing all routes that are received and accepted from the specified neighbor:

```
switch# show bgp ipv4 unicast neighbors 9.0.0.2 routes
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete

VRF : default
Local Router-ID 9.0.0.1

    Network            Nexthop         Metric     LocPrf     Weight Path
--------------------------------------------------------------------------
*>e 9.0.0.0/24         9.0.0.2         0          100        0      200 65534.65535
3.4 18.54934 3574.8570 5.6 ?
*>e 100.0.0.0/24       9.0.0.2         0          100        0      200 ?
*>e 100.0.1.0/24       9.0.0.2         0          100        0      200 ?
*>e 100.0.2.0/24       9.0.0.2         0          100        0      200 ?
*>e 100.0.3.0/24       9.0.0.2         0          100        0      200 ?
*ae 100.0.3.0/24       9.0.0.3         0          100        0      200 ?
Total number of entries 6
```

Showing 12 VPN EVPN routes that are received and accepted from the specified neighbor:

```
switch# show bgp l2vpn evpn neighbor 9.0.0.2 routes

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
Origin codes: i - IGP, e - EGP, ? - incomplete
EVPN Route-Type 2 prefix: [2]:[ESI]:[EthTag]:[MAC]:[OrigIP]
EVPN Route-Type 3 prefix: [3]:[EthTag]:[OrigIP]

VRF            : default
Local Router-id: 9.0.0.1

   Network            Nexthop         Metric     LocPrf     Weight Path
```

```
------------------------------------------------------------------------
*>e 9.0.0.0/24          9.0.0.2         0           0           0      200 65534.65535
3.4 18.54934 3574.8570 5.6 ?
*>e 100.0.0.0/24        9.0.0.2         0           0           0      200 ?
*>e 100.0.1.0/24        9.0.0.2         0           0           0      200 ?
*>e 100.0.2.0/24        9.0.0.2         0           0           0      200 ?
*>e 100.0.3.0/24        9.0.0.2         0           0           0      200 ?
*ae 100.0.3.0/24        9.0.0.3         0           100         0      200 ?
Total number of entries 6
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show bgp neighbors

```
show bgp [vrf <VRF-NAME>] {ipv4 unicast | ipv6 unicast | all}
     neighbors [vsx-peer]
show bgp l2vpn evpn neighbors
```

Description

This command shows information about BGP and TCP connections to neighbors. If neighbors are member of a peer-group, the command shows the configured values inherited from the peer-group. The configured values are postfixed with a caret (^) for inherited values.

Parameter	Description
`vrf <VRF-NAME>`	Shows the information for a specified VRF.
`ipv4 unicast`	Shows the information for an IPv4 unicast address family.
`ipv6 unicast`	Shows the information for an IPv6 unicast address family.
`all`	Shows the information for all address families and subaddress families.

Parameter	Description
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.
`l2vpn evpn`	Shows the information for L2VPN EVPN address family.

Examples

Showing all information about BGP and TCP connections to neighbors:

```
switch# show bgp all neighbors
Codes: ^ Inherited from peer-group
VRF : default

BGP Neighbor 10.1.1.2 (Internal)
    Description         :
    Peer-group          :

    Remote Router Id    : 10.1.1.2          Local Router Id    : 1.0.0.1
    Remote AS           : 1                 Local AS           : 1
    Remote Port         : 0                 Local Port         : 0
    State               : Established       Admin Status       : Up
    Conn. Established   : 1                 Conn. Dropped      : 0
    Passive             : No                Update-Source      :
    Cfg. Hold Time      : 180               Cfg. Keep Alive    : 60
    Neg. Hold Time      : 0                 Neg. Keep Alive    : 0
    Up/Down Time        : 00h:00m:00s       Connect-Retry Time : 120
    Local-AS Prepend    : No                Alt. Local-AS      : 0
    BFD                 : Disabled          Slow Peer          : No
    Password            :
    Last Err Sent       : No Error
    Last SubErr Sent    : No Error
    Last Err Rcvd       : No Error
    Last SubErr Rcvd    : No Error

    Graceful-Restart    : Enabled           Gr. Restart Time   : 120
    Gr. Stalepath Time  : 150               Remove Private-AS  : No
    TTL                 : 255               Local Cluster-ID   :
    Weight              : 0                 Fall-over          : No

    Message statistics       Sent     Rcvd
    ------------------       -----    -----
    Open                        7        6
    Notification                5        2
    Updates                     3        2
    Keepalives                 12       10
    Route Refresh               0        0
    Total                      28       20

    Capability                          Advertised      Received
    ----------                          ----------      ----------
    Route Refresh                       Yes             No
    Graceful Restart                    Yes             No
    Four Octet ASN                      Yes             No
    Add-Path                            Yes             Yes
    Address family IPv4 Unicast         Yes             No
    Address family IPv6 Unicast         No              No
    Address family L2VPN EVPN           No              No
```

```
    Address Family : IPv4 Unicast
    ------------------------------
    Rt. Reflect. Client : No                  Send Community    :
    Allow-AS in         : 0                   Advt. Interval    : 30
    Max. Prefix         : 64000               Soft Reconfig In  :
    Nexthop-Self        :                     Default-Originate :
    Update-Group        : 2

    Cfg. Add-Path       : Send and Receive
    Neg. Add-Path       : Send and Receive
    Routemap In         :
    Routemap Out        :
    ORF type            : Prefix-list
    ORF capability      : Receive

    Address Family : IPv6 Unicast
    ------------------------------
    Rt. Reflect. Client : No                  Send Community    :
    Allow-AS in         : 0                   Advt. Interval    : 30
    Max. Prefix         : 64000               Soft Reconfig In  :
    Nexthop-Self        :                     Default-Originate :
    Update-Group        : 3

    Cfg. Add-Path       : Send and Receive
    Neg. Add-Path       : Send and Receive
    Routemap In         :
    Routemap Out        :
    ORF type            : Prefix-list
    ORF capability      : Receive

    Address Family : L2VPN EVPN
    ------------------------------
    Rt. Reflect. Client : No                  Send Community    : extended
    Allow-AS in         : 0                   Advt. Interval    : 30
    Max. Prefix         : 32768               Soft Reconfig In  :
    Nexthop-Self        :                     Default-Originate :
    Update-Group        : 4

    Cfg. Add-Path       :
    Neg. Add-Path       :
    Routemap In         :
    Routemap Out        :
    ORF type            : Prefix-list
    ORF capability      : Receive
```

```
Showing information about L2VPN EVPN connections to neighbors:
                switch# show bgp l2vpn evpn neighbors
Codes: ^ Inherited from peer-group, * Dynamic Neighbor

VRF : default

BGP Neighbor 10.1.1.2 (Internal)
    Description         :
    Peer-group          :

    Remote Router Id    : 10.1.1.2          Local Router Id    : 10.1.1.1
    Remote AS           : 1                 Local AS           : 1
    Remote Port         : 179               Local Port         : 56008
    State               : Established       Admin Status       : Up
    Conn. Established   : 1                 Conn. Dropped      : 0
    Passive             : No                Update-Source      :
```

```
    Cfg. Hold Time      : 180               Cfg. Keep Alive    : 60
    Neg. Hold Time      : 180               Neg. Keep Alive    : 60
    Up/Down Time        : 00m:01w:03d       Alt. Local-AS      : 0
    Local-AS Prepend    : No
    BFD                 : Disabled
    Password            :
    Last Err Sent       : No Error
    Last SubErr Sent    : No Error
    Last Err Rcvd       : No Error
    Last SubErr Rcvd    : No Error

    Graceful-Restart    : Enabled           Gr. Restart Time   : 120
    Gr. Stalepath Time  : 150               Remove Private-AS  : No
    TTL                 : 255               Local Cluster-ID   :
    Weight              : 0                 Fall-over          : No

    Message statistics       Sent     Rcvd
    ------------------       -----    -----
    Open                        1        1
    Notification                0        0
    Updates                     3        2
    Keepalives              17995    18009
    Route Refresh               0        0
    Total                   17999    18012

    Capability                          Advertised      Received
    ----------                          ----------      ----------
    Route Refresh                       Yes             Yes
    Graceful Restart                    Yes             Yes
    Four Octet ASN                      Yes             Yes
    Address family IPv4 Unicast         Yes             Yes
    Address family IPv6 Unicast         Yes             Yes
    Address family L2VPN EVPN           Yes             Yes

    Address Family : L2VPN EVPN
    ---------------------------

    Rt. Reflect. Client : No                Send Community    : extended
    Allow-AS in         : 0                 Advt. Interval    : 30
    Max. Prefix         : 32768             Soft Reconfig In  :
    Nexthop-Self        :                   Default-Originate :

    Routemap In         :
    Routemap Out        :
```

Showing information for BGP IPv4 unicast neighbors:

```
switch# show bgp ipv4 unicast neighbors
Codes: ^ Inherited from peer-group

VRF : default

BGP Neighbor 10.1.1.2 (Internal)
    Description         :
    Peer-group          :

    Remote Router Id    : 10.1.1.2          Local Router Id    : 1.0.0.1
    Remote AS           : 1                 Local AS           : 1
    Remote Port         : 0                 Local Port         : 0
    State               : Idle              Admin Status       : Up
    Conn. Established   : 0                 Conn. Dropped      : 0
```

```
Passive             : No                Update-Source      :
Cfg. Hold Time      : 180               Cfg. Keep Alive    : 60
Neg. Hold Time      : 0                 Neg. Keep Alive    : 0
Up/Down Time        : 00h:00m:00s       Connect-Retry Time : 120
Local-AS Prepend    : No                Alt. Local-AS      : 0
BFD                 : Disabled          Slow Peer          : Yes
Password            :
Last Err Sent       : No Error
Last SubErr Sent    : No Error
Last Err Rcvd       : No Error
Last SubErr Rcvd    : No Error

Graceful-Restart    : Enabled           Gr. Restart Time   : 120
Gr. Stalepath Time  : 150               Remove Private-AS  : No
TTL                 : 255               Local Cluster-ID   :
Weight              : 0                 Fall-over          : No

Message statistics       Sent     Rcvd
------------------       -----    -----
Open                     0        0
Notification             0        0
Updates                  0        0
Keepalives               0        0
Route Refresh            0        0
Total                    0        0

Capability                          Advertised      Received
----------                          ----------      ----------
Route Refresh                       Yes             No
Graceful Restart                    Yes             No
Four Octet ASN                      Yes             No
Add-Path                            Yes             Yes
Address family IPv4 Unicast         Yes             No
Address family IPv6 Unicast         No              No
Address family L2VPN EVPN           Yes             Yes

Address Family : IPv4 Unicast
-----------------------------

Rt. Reflect. Client : No                Send Community    :
Allow-AS in         : 0                 Advt. Interval    : 30
Max. Prefix         : 64000             Soft Reconfig In  :
Nexthop-Self        :                   Default-Originate :
Update-Group        : 12

Routemap In         :
Routemap Out        :

ORF Type            : Prefix-list
ORF Capability      : Send

Cfg. Add-Path       : Send
Neg. Add-Path       : Send
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show bgp paths

```
show bgp [vrf <VRF-NAME>] {ipv4 unicast | ipv6 unicast | all}
     paths [vsx-peer]
show bgp l2vpn evpn paths
```

Description

Shows received BGP path information in the database.

Parameter	Description
`vrf <VRF-NAME>`	Shows the information for a specified VRF.
`ipv4 unicast`	Shows the information for an IPv4 unicast address family.
`ipv6 unicast`	Shows the information for an IPv6 unicast address family.
`all`	Shows the information for all address families and subaddress families.
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.
`l2vpn evpn`	Shows the information for L2VPN EVPN address family.

Examples

Showing received BGP path information from the specified IPv4 unicast neighbor:

```
switch# show bgp ipv4 unicast paths
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths
VRF : default
Local Router-ID 9.0.0.1

    Network            Nexthop          PathID      Path
-----------------------------------------------------------
*>e 9.0.0.0/24         9.0.0.2          0           200 65534.65535 3.4 18.54934
3574.8570 5.6
*>e 100.0.0.0/24       9.0.0.2          0           200
*>e 100.0.1.0/24       9.0.0.2          0           200
```

```
*>e 100.0.2.0/24       9.0.0.2          0           200
*>e 100.0.3.0/24       9.0.0.2          10          200
*ae  100.0.3.0/24      9.0.0.3          5           200
Total number of entries 6
```

Showing received BGP path information from the specified L2VPN EVPN neighbor:

```
switch# show bgp l2vpn evpn paths

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete
EVPN Route-Type 2 prefix: [2]:[ESI]:[EthTag]:[MAC]:[OrigIP]
EVPN Route-Type 3 prefix: [3]:[EthTag]:[OrigIP]

VRF : default
Local Router-ID 9.0.0.2

     Network                                                Nexthop         Path
-------------------------------------------------------------------------------
Route Distinguisher: 10.1.1.54:32967    (L2VNI 30000)
*>   [2]:[0]:[0]:[00:06:f6:3f:e3:c1]:[]                    1.1.1.20          100
*>   [2]:[0]:[0]:[8c:60:4f:f2:f5:41]:[]                    1.1.1.10          100
Total number of entries 2
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show bgp peer-group summary

```
show bgp [vrf <VRF-NAME>] {ipv4 unicast |
     ipv6 unicast | all} peer-group <PEER-GROUP-NAME>
     summary [vsx-peer]
show bgp l2vpn evpn peer-group <PEER-GROUP-NAME> summary
```

Description

This command shows the peer-group information in the database.

Parameter	Description
`vrf <VRF-NAME>`	Shows the information for a specified VRF.
`ipv4 unicast`	Shows the information for an IPv4 unicast address family.
`ipv6 unicast`	Shows the information for an IPv6 unicast address family.
`all`	Shows the information for all address families and subaddress families.
<PEER-GROUP-NAME>	Shows the information for the BGP peer-group for the BGP instance.
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.
`l2vpn evpn`	Shows the information for L2VPN EVPN address family. This parameter only applies to HPE Aruba Networking 8100, 8325, 9300/9300S, 8360 Switch series.

Examples

Showing the information from IPv4 unicast address families in `pg_name1` peer-group:

```
switch# show bgp ipv4 unicast peer-group pg_name1 summary
Codes: * Dynamic Neighbor
VRF : default
BGP Peer-Group Summary
======================
 Local AS               : 1           BGP Router Identifier  : 2.2.2.2
 Peers                  : 1           Dynamic Peer Count     : 3
 Cfg. Hold Time         : 180         Cfg. Keep Alive        : 60

 Neighbor        Remote-AS MsgRcvd MsgSent Up/Down Time State       AdminStatus
 10.0.0.1        1         8       10      00h:00m:58s  Established Up
*10.1.1.5        11        15      14      00h:10m:24s  Established Up
```

Showing the information from all address families in `pg_name1` peer-group:

```
switch# show bgp all unicast peer-group pg_name1 summary
Codes: * Dynamic Neighbor
VRF : default
BGP Peer-Group Summary
======================
 Local AS               : 1           BGP Router Identifier  : 2.2.2.2
 Peers                  : 1           Dynamic Peer Count     : 3
 Cfg. Hold Time         : 180         Cfg. Keep Alive        : 60
 Confederation Id       : 0
For address family: IPv4 Unicast
 Neighbor        Remote-AS MsgRcvd MsgSent Up/Down Time State       AdminStatus
 10.0.0.1        1         8       10      00h:00m:58s  Established Up
*10.1.1.5        11        15      14      00h:10m:24s  Established Up
For address family: IPv6 Unicast
 Neighbor        Remote-AS MsgRcvd MsgSent Up/Down Time State       AdminStatus
  1001::1002     11        12      12      00h:00m:07s  Established Up
  2001::2002     11        12      12      00h:00m:07s  Established Up
```

```
For address family: L2VPN EVPN
 Neighbor        Remote-AS MsgRcvd MsgSent Up/Down Time State       AdminStatus
  10.0.0.1        1         8       10      00h:00m:58s  Established Up
  10.1.1.6        11        15      14      00h:10m:24s  Established Up
```

Showing the information from L2VPN EVPN address families in `pg_name1` peer-group:

```
switch# show bgp l2vpn evpn peer-group pg_name1 summary
VRF : default
BGP Peer-Group Summary
======================
 Local AS               : 1           BGP Router Identifier  : 2.2.2.2
 Peers                  : 1           Dynamic Peer Count     : 3
 Cfg. Hold Time         : 180         Cfg. Keep Alive        : 60
 Confederation Id       : 0


 Neighbor        Remote-AS MsgRcvd MsgSent Up/Down Time State       AdminStatus
 10.0.0.1        1         8       10      00h:00m:58s  Established Up
*10.1.1.6        11        15      14      00h:10m:24s  Established Up
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show bgp summary

```
show bgp [vrf <VRF-NAME>] {ipv4 unicast | ipv6 unicast | all}
     summary [vsx-peer]
show bgp l2vpn evpn summary
```

Description

This command shows a summary of the status of Border Gateway Protocol (BGP) connections.

Parameter	Description
`ipv4 unicast`	Selects to display the BGP summary information for the IPv4 subaddress family identifier.
`ipv6 unicast`	Selects to display the BGP summary information for the IPv6 subaddress family identifier.
`all`	Selects to display the BGP summary information for all VRFs and address-families.
`vrf <VRF-NAME>`	Selects to display information by VRFs by specifying the VRF name.
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.
`l2vpn evpn`	Shows the information for L2VPN EVPN address family. This parameter only applies to HPE Aruba Networking 8100, 8325, 8360, 9300, 9300S Switch series

Examples

Showing BGP summary information for all address-families:

```
switch(config-bgp)# show bgp all summary
Codes: * Dynamic Neighbor
VRF : default
BGP Summary
 Local AS               : 100          BGP Router Identifier  : 9.0.0.1
 Peers                  : 1            Log Neighbor Changes   : No
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60
 Confederation Id       : 0

Address-family : IPv4 Unicast
------------------------------
 Neighbor        Remote-AS MsgRcvd MsgSent Up/Down Time State       AdminStatus
 9.0.0.2         200       25      23      00h:17m:50s  Established Up
*10.1.1.5        11        26      24      00h:20m:26s  Established Up


Address-family : IPv6 Unicast
------------------------------
 Neighbor        Remote-AS MsgRcvd MsgSent Up/Down Time State       AdminStatus
*2001::2002      11        3       3       00h:00m:14s  Established Up
 9000::2         200       25      23      00h:17m:50s  Established Up


Address-family : VPNv4 Unicast
------------------------------
 Neighbor        Remote-AS MsgRcvd MsgSent Up/Down Time State       AdminStatus
 1.1.1.1         100       207     208     02h:54m:18s  Established Up
*3.3.3.4         11        26      24      00h:20m:26s  Established Up

Address-family : L2VPN EVPN
------------------------------
 Neighbor        Remote-AS MsgRcvd MsgSent Up/Down Time State       AdminStatus
 10.0.0.2        200       25      23      00h:17m:50s  Established Up
```

```
*10.1.1.6        11        26      24      00h:20m:26s  Established Up
VRF : v1
BGP Summary
 Local AS               : 100          BGP Router Identifier  : 9.0.0.1
 Peers                  : 1            Log Neighbor Changes   : No
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60

Address-family : IPv4 Unicast
------------------------------
 Neighbor        Remote-AS MsgRcvd MsgSent Up/Down Time State       AdminStatus
*4.4.4.4         11        26      24      00h:20m:26s  Established Up
 9.0.0.2         200       25      23      00h:17m:50s  Established Up

Address-family : IPv6 Unicast
------------------------------
 Neighbor        Remote-AS MsgRcvd MsgSent Up/Down Time State       AdminStatus
*3001::3002      11        3       3       00h:00m:14s  Established Up
 9000::2         200       25      23      00h:17m:50s  Established Up
```

Showing BGP summary information for a specific VRF for IPv4 unicast network:

```
switch(config-bgp)# show bgp ipv4 unicast vrf v1 summary
Codes: * Dynamic Neighbor
VRF : v1
BGP Summary
 Local AS               : 100          BGP Router Identifier  : 9.0.0.1
 Peers                  : 1            Log Neighbor Changes   : No
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60

 Neighbor        Remote-AS MsgRcvd MsgSent Up/Down Time State       AdminStatus
 9.0.0.2         200       25      23      00h:17m:50s  Established Up
*10.1.1.5        11        26      24      00h:20m:26s  Established Up
```

Showing BGP summary information for L2VPN EVPN:

```
switch(config-bgp)# do show bgp l2vpn evpn summary
Codes: * Dynamic Neighbor
VRF : default
BGP Summary
 Local AS               : 100          BGP Router Identifier  : 9.0.0.1
 Peers                  : 1            Log Neighbor Changes   : No
 Cfg. Hold Time         : 180          Cfg. Keep Alive        : 60

 Neighbor        Remote-AS MsgRcvd MsgSent Up/Down Time State       AdminStatus
 10.0.0.2        200       25      23      00h:17m:50s  Established Up
 10.1.1.6        11        26      24      00h:20m:26s  Established Up
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show bgp l2vpn evpn

```
show bgp l2vpn evpn <ROUTE-TYPE>
```

Description

Shows the BGP L2VPN information.

Parameter	Description
<ROUTE-TYPE>	Specifies the routes filtered by NLRI route type. HPE Aruba Networking 8325 and 10000 Switch series can display information for route types 1-6. HPE Aruba Networking 6300, 6400, 8100, 8320, 8360, 9300, 9300S Switch series display information for route types 2-6.

Examples

Showing BGP L2VPN information for extcommunity:

```
switch# show bgp l2vpn evpn extcommunity
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

EVPN Route-Type 1 prefix: [1]:[ESI]:[EthTag]
EVPN Route-Type 2 prefix: [2]:{ESI]:{EthTag]:[MAC]:[OrigIP]
EVPN Route-Type 3 prefix: [3]:[EthTag]:[OrigIP]
EVPN Route-Type 4 prefic: [4]:[ESI]:[OrigIP]
EVPN Route-Type 5 prefix: [5]:[ESI]:[EthTag]:[IPAddrLen]:[IPAddr]
EVPN Route-Type 6 prefix: [6]:[EthTag}:McastSrcIP]:[McastGrpIP]:[OrigIP]


VRF : default
Local Router-ID 2.2.2.2


Network                                                 Nexthop
       Ext-Community
-------------------------------------------------------------------------------
-------
Route Distinguisher: 2:2                (L2VNI 1)
*>i [2]:[0]:[0]:[00:00:00:00:00:aa]:[100.100.1.2]  1.1.1.1
       1:1,100:100,00:00:00:00:00:11
*>i [3]:[0]:[1.1.1.1]                              1.1.1.1
```

```
        1:1

Route Distinguisher: 2:2                (L2VNI 2)
*>i [2]:[0]:[0]:[00:00:00:00:00:aa]:[200.200.1.2]  3.3.3.3
       2:2,100:100,00:00:00:00:00:22
*>i [3]:[0]:[3.3.3.3]                              3.3.3.3
        2:2

Route Distinguisher: 1:1                (L3VNI 10000)
*>i [2]:[0]:[0]:[00:00:00:00:00:aa]:[100.100.1.2]  1.1.1.1
        1:1,100:100,00:00:00:00:00:11

Route Distinguisher: 2:2                (L3VNI 10000)
*>i [2]:[0]:[0]:[00:00:00:00:00:aa]:[200.200.1.2]  3.3.3.3
        2:2,100:100,00:00:00:00:00:22

Route Distinguisher: 4.4.4.4:40         (L2VNI 400)
*>i [3]:[0]:[4.4.4.4]                     4.4.4.4
        4:4,IGMP-MLD Proxy

Total number of entries 7
```

Showing BGP L2VPN information:

```
switch# show bgp l2vpn evpn 5.5.5.5:41-[6]:[0]:[30.1.36.30]:[232.0.0.8:
VRF : default
BGP Local AS 65000       BGP Router-id 3.3.3.3

Network                   : 5.5.5.5:41-[6]:[0]:[30.1.36.30]:[232.0.0.8]:[5.5.5.5]
Nexthop                   : 5.5.5.5
vni                       : 0                    vni_type          : L2VNI
Peer                      : 2.2.2.2              Origin            : incomplete
Metric                    : 0                    Local Pref        : 100
Weight                    : 0                    Calc. Local Pref  : 100
Best                      : Yes                  Valid             : Yes
Type                      : internal             Stale             : No
Originator ID             : 5.5.5.5
Aggregator ID             :
Aggregator AS             :
Atomic Aggregate          :

AS-Path                   :
Cluster List              : 2.2.2.2
Communities               :
Ext-Communities           : RT: 139:139
```

Showing BGP L2VPN information for route type 4:

```
switch# show bgp l2vpn evpn route-type 4

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

EVPN Route-Type 1 prefix: [1]:[ESI]:[EthTag]
EVPN Route-Type 2 prefix: [2]:[ESI]:[EthTag]:[MAC]:[OrigIP]
EVPN Route-Type 3 prefix: [3]:[EthTag]:[OrigIP]
EVPN Route-Type 4 prefix: [4]:[ESI]:[OrigIP]
EVPN Route-Type 5 prefix: [5]:[ESI]:[EthTag]:[IPAddrLen]:[IPAddr]
```

```
VRF : default
Local Router-ID 16.16.16.16

Network                                          Nexthop  Metric  LocPrf  Weight
Path
-------------------------------------------------------------------------------
--
Route Distinguisher: 2.2.2.2:840          (L2VNI 0)
*>i [4]:[0000.0000.0001.0000.0001]:[2.2.2.2]    2.2.2.2  0       0       0      i
Route Distinguisher: 2.2.2.2:860          (L2VNI 0)
*>i [4]:[0000.0000.0001.0000.0002]:[2.2.2.2]    2.2.2.2  0       0       0      i
Route Distinguisher: 3.3.3.3:800          (L2VNI 0)
*>i [4]:[0000.0000.0001.0000.0001]:[3.3.3.3]    3.3.3.3  0       0       0      i
Route Distinguisher: 3.3.3.3:820          (L2VNI 0)
*>i [4]:[0000.0000.0001.0000.0003]:[3.3.3.3]    3.3.3.3  0       0       0      i
Total number of entries 4
```

Showing BGP L2VPN information for route type 6:

```
switch# show bgp l2vpn evpn route-type 6

Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete
EVPN Route-Type 6 prefix: [6]:[EthTag]:[McastSrcIP]:[McastGrpIP]:[OrigIP]

VRF: default
Router-ID not configured

Network                                         Nexthop   Metric LocPrf  Weight
Path
-------------------------------------------------------------------------------
--
Route Distinguisher: 10.1.1.54:32967
*>  [6]:[0]:[0.0.0.0]:[239.1.1.1]:[1.1.1.1]    1.1.1.1   0      100     32768   i
*>  [6]:[0]:[10.1.1.30]:[225.1.1.1]:[1.1.1.1]  1.1.1.1   0      100     0       i
*>  [6]:[0]:[10.1.1.40]:[225.1.1.2]:[1.1.1.1]  1.1.1.1   0      100     0       ?
*>  [6]:[0]:[0.0.0.0]:[0.0.0.0]:[10.1.1.1]     10.1.1.1  0      100     0       i

Route Distinguisher: 2:2
*>  [6]:[0]:[10.1.1.30]:[232.1.1.1]:[1.1.1.1]  1.1.1.1   0      100     0       ?
*>  [6]:[0]:[10.1.1.40]:[239.1.1.1]:[2.2.2.2]  2.2.2.2   0      100             0
Total number of entries 5
```

Command History

Release	Modification
10.17	Expanded route type range to include route type 7.
10.16	Expanded route type range to include route type 6.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

show bgp l2vpn evpn vni route-type

`show bgp l2vpn evpn vni <VNI-Value> route-type <ROUTE-TYPE-Value>`

Description

Shows the BGP L2VPN information for the particular EVPN VNI and routes type.

Parameter	Description
<VNI-Value>	Specifies the VNI.
<ROUTE-TYPE-Value>	Specifies the routes filtered by NLRI route type. HPE Aruba Networking 8325 and 10000 Switch series can display information for route types 1-6. HPE Aruba Networking 6300, 6400, 8100, 8320, 8360, 9300, 9300S Switch series display information for route types 2-6.

Examples

Showing BGP L2VPN information for the particular EVPN VNI and route type:

```
switch# show bgp l2vpn evpn vni 30000 route-type 5
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete
EVPN Route-Type 5 prefix: [5]:[ESI]:[EthTag]:[IPAddrLen]:[IPAddr]

VRF : default
Router-ID not configured
Network                                          Nexthop         Metric LocPrf
Weight Path
-------------------------------------------------------------------------------
---------
Route Distinguisher: 1:100              (L3VNI 10000)
*>i [5]:[0]:[0]:[24]:[32.32.32.0]                3.3.3.3         0      100    0
  ?
*>  [5]:[0]:[0]:[24]:[52.52.52.0]                1.1.1.1         0      100    0
  ?
*>i [5]:[0]:[0]:[64]:[aaa::]                     3.3.3.3         0      100    0
  ?
```

```
Total number of entries 3
```

Command History

Release	Modification
10.16	Expanded route type range to include route type 6.
10.14.1000	This command can display information for route type **1** on HPE Aruba Networking 8325 and 10000 Switch series
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

show bgp l2vpn evpn vtep

`show bgp l2vpn evpn vtep <IP-address>`

Description

Shows the BGP L2VPN information for the particular EVPN VTEP IP address.

Parameter	Description
<IP-address>	Specifies the VTEP IP address.

Examples

Showing BGP L2VPN information for the particular EVPN VTEP IP:

```
switch# show bgp l2vpn evpn vtep 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

EVPN Route-Type 2 prefix: [2]:[ESI]:[EthTag]:[MAC]:[OrigIP]
EVPN Route-Type 3 prefix: [3]:[EthTag]:[OrigIP]
EVPN Route-Type 5 prefix: [5]:[ESI]:[EthTag]:[IPAddrLen]:[IPAddr]
```

```
VRF : default
Local Router-ID 2.2.2.2

     Network                                                Nexthop          Metric
   LocPrf    Weight   Path
-------------------------------------------------------------------------------
---------------------------
Route Distinguisher: 1.1.1.1:2            (L2VNI 2)
*>i [2]:[0]:[0]:[00:00:00:00:00:33]:[10.1.1.10]           1.1.1.1          0
   100        0       ?
*>i [2]:[0]:[0]:[00:00:00:00:00:33]:[1000::10]            1.1.1.1          0
   100        0       ?
*>i [2]:[0]:[0]:[00:50:56:96:15:1c]:[10.1.1.1]            1.1.1.1          0
   100        0       ?
*>i [2]:[0]:[0]:[00:50:56:96:15:1c]:[]                    1.1.1.1          0
   100        0       ?
*>i [3]:[0]:[1.1.1.1]                                     1.1.1.1          0
   100        0       ?

Route Distinguisher: 1.1.1.1:2            (L3VNI 10000)
*>i [2]:[0]:[0]:[00:00:00:00:00:33]:[10.1.1.10]           1.1.1.1          0
   100        0       ?
*>i [2]:[0]:[0]:[00:00:00:00:00:33]:[1000::10]            1.1.1.1          0
   100        0       ?
*>i [2]:[0]:[0]:[00:50:56:96:15:1c]:[10.1.1.1]            1.1.1.1          0
   100        0       ?
*>i [2]:[0]:[0]:[00:50:56:96:15:1c]:[]                    1.1.1.1          0
   100        0       ?

Route Distinguisher: 1:100                (L3VNI 10000)
*>i [5]:[0]:[0]:[24]:[10.1.1.0]                           1.1.1.1          0
   100        0       ?
*>i [5]:[0]:[0]:[64]:[1000::]                             1.1.1.1          0
   100        0       ?
Total number of entries 11
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

show bgp l2vpn evpn vtep route-type

```
show bgp l2vpn evpn vtep <IP-address> route-type <ROUTE-TYPE-Value>
```

Description

Shows the BGP L2VPN information for the particular EVPN VTEP IP address and routes type.

Parameter	Description
<IP-address>	Specifies the VTEP IP address.
<ROUTE-TYPE-Value>	Specifies the routes filtered by NLRI route type. HPE Aruba Networking 8325 and 10000 Switch series can display information for route types 1-5. HPE Aruba Networking 6300, 6400, 8100, 8320, 8360, 9300, 9300S switches display information for route types 2, 3 and 5.

Examples

Showing BGP L2VPN information for the particular EVPN VTEP and route type:

```
switch# show bgp l2vpn evpn vtep 1.1.1.1 route-type 5
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete
EVPN Route-Type 5 prefix: [5]:[ESI]:[EthTag]:[IPAddrLen]:[IPAddr]

VRF : default
Router-ID not configured

     Network                                                    Nexthop
Metric LocPrf Weight Path
-------------------------------------------------------------------------------
----------------------
Route Distinguisher: 1:100              (L3VNI 10000)
*>i [5]:[0]:[0]:[24]:[32.32.32.0]                              1.1.1.1         0
   100    0      ?
*>  [5]:[0]:[0]:[24]:[52.52.52.0]                              1.1.1.1         0
   100    0      ?
*>i [5]:[0]:[0]:[64]:[aaa::]                                   1.1.1.1         0
   100    0      ?

Total number of entries 3
```

Command History

Release	Modification
10.14.1000	This command can display information for route type **1** on 8325 and 10000 Switch series.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

show bgp l2vpn evpn vtep vni

`show bgp l2vpn evpn vtep <IP-address> vni <VNI-Value>`

Description

Shows the BGP L2VPN information for the particular EVPN VTEP IP address and VNI.

Parameter	Description
<IP-address>	Specifies the VTEP IP address.
<VNI-Value>	Specifies the VNI.

Examples

Showing BGP L2VPN information for the particular EVPN VTEP IP and VNI:

```
switch# show bgp l2vpn evpn vtep 1.1.1.1 vni 10000
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

EVPN Route-Type 2 prefix: [2]:[ESI]:[EthTag]:[MAC]:[OrigIP]
EVPN Route-Type 3 prefix: [3]:[EthTag]:[OrigIP]
EVPN Route-Type 5 prefix: [5]:[ESI]:[EthTag]:[IPAddrLen]:[IPAddr]
VRF : default
Local Router-ID 2.2.2.2

     Network                                                Nexthop         Metric
   LocPrf    Weight   Path
-------------------------------------------------------------------------------
-------------------------
Route Distinguisher: 1.1.1.1:2            (L3VNI 10000)
*>i [2]:[0]:[0]:[00:00:00:00:00:33]:[10.1.1.10]           1.1.1.1         0
   100        0       ?
*>i [2]:[0]:[0]:[00:00:00:00:00:33]:[1000::10]            1.1.1.1         0
   100        0       ?
*>i [2]:[0]:[0]:[00:50:56:96:15:1c]:[10.1.1.1]            1.1.1.1         0
   100        0       ?
*>i [2]:[0]:[0]:[00:50:56:96:15:1c]:[]                    1.1.1.1         0
   100        0       ?
```

```
Route Distinguisher: 1:100                (L3VNI 10000)
*>i [5]:[0]:[0]:[24]:[10.1.1.0]                              1.1.1.1         0
   100        0       ?
*>i [5]:[0]:[0]:[64]:[1000::]                                1.1.1.1         0
   100        0       ?
Total number of entries 6
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

show bgp l2vpn evpn vtep vni route-type

`show bgp l2vpn evpn vtep <VTEP-ID> vni <VNI-Value> route-type <ROUTE-TYPE-Value>`

Description

Shows the BGP L2VPN information for the particular EVPN VTEP, VNI, and route type.

Parameter	Description
<VTEP-ID>	Specifies the VTEP.
<VNI-Value>	Specifies the VNI.
<ROUTE-TYPE-Value>	Specifies the route type. HPE Aruba Networking 8325 and 10000 Switch series can display information for route types 1-5. HPE Aruba Networking 6300, 6400, 8100, 8320, 8360, 9300/9300S Switch series display information for route types 2-5.

Examples

Showing BGP L2VPN information for the particular EVPN VTEP, VNI and route type 2:

```
switch# show bgp l2vpn evpn vtep 1.1.1.1 vni 10000 route-type 2
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

EVPN Route-Type 2 prefix: [2]:[ESI]:[EthTag]:[MAC]:[OrigIP]
EVPN Route-Type 3 prefix: [3]:[EthTag]:[OrigIP]
EVPN Route-Type 5 prefix: [5]:[ESI]:[EthTag]:[IPAddrLen]:[IPAddr]
VRF : default
Local Router-ID 2.2.2.2

     Network                                          Nexthop    Metric  LocPrf
Weight   Path
-------------------------------------------------------------------------------
---------
Route Distinguisher: 1.1.1.1:2   (L3VNI 10000)
*>i [2]:[0]:[0]:[00:50:56:96:7d:03]:[10.1.1.1]    1.1.1.1    0       100        0
     ?
*>i [2]:[0]:[0]:[00:50:56:96:7d:03]:[]            1.1.1.1    0       100        0
     ?

Total number of entries 3
```

Showing BGP L2VPN information for the particular EVPN VTEP, VNI and route type 1, and VNI on a 10000 or 8325 Switch series:

```
switch# show bgp l2vpn evpn vni 30000 route-type 1
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete
EVPN Route-Type 1 prefix: [1]:[ESI]:[EthTag]
EVPN Route-Type 2 prefix: [2]:[ESI]:[EthTag]:[MAC]:[OrigIP]
EVPN Route-Type 3 prefix: [3]:[EthTag]:[OrigIP]
EVPN Route-Type 5 prefix: [5]:[ESI]:[EthTag]:[IPAddrLen]:[IPAddr]
VRF : default
Router-ID not configured
Network                                Nexthop         Metric LocPr  Weight Path
-------------------------------------------------------------------------
Route Distinguisher: 2:2 (L2VNI 30000)
*>   [1]:[0000.0000.0010.0000][0]     1.1.1.1         0       100    32768  i
Total number of entries 1
```

Command History

Release	Modification
10.14.1000	When issued from an HPE Aruba Networking 8325 or 10000 Switch series, this command can display information for route type **1**.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except	Operator (>) or Manager	Administrators or local user group members with execution

Platforms	Command context	Authority
for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	(#)	rights for this command.

show bgp unicast protected

```
show bgp ip unicast protected
```

Description

Displays BGP routes with options to view specific protected IPs or routes when BGP-PIC is configured.

Examples

Showing BGP IP routes:

```
9300# show bgp ip unicast protected
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, e external S Stale, R Removed, a additional-paths,
              b backup path
Origin codes: i - IGP, e - EGP, ? - incomplete
VRF : default
Local Router-ID 200.0.0.1
Network           Nexthop                                  Metric     LocPrf
Weight Path
*>i 201.0.0.0/24                                  20.20.20.2
     0          100        0       i
*bi 201.0.0.0/24                                  40.40.40.2
     0          100        0       i
Total number of entries 2
```

Command History

Release	Modification
10.14.1000	Command introduced.

Command Information

Platforms	Command context	Authority
9300	Operator (>) or Manager (#)	

show ip ospf routes

```
show ip ospf [<PROCESS-ID>] routes
     [<IPV4-ADDR>/<MASK>] protected [all-vrfs | vrf <VRF-NAME|
```

Description

Displays OSPF routing information.

Parameter	Description
<IPV4-ADDR>	Specify an IP address in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.
<MASK>	Specifies the number of bits in the address mask in CIDR format (**x**), where **x** is a decimal number from 0 to 32.
<PROCESS-ID>	Enter an OSPF process ID to display OSPF neighbor information for the particular OSPF process. Range: 1 to 65535.
`all-vrfs`	Select to display neighbor information for all VRFs.
`vrf <VRF-NAME>`	Specify the name of a VRF. Default: default.
`protected`	Selects a protected route. Specifying an IPv4 address will select protected route for that prefix.

Examples

Showing OSPF routing table information:

```
switch# show ip ospf routes
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF default, Routing Table
--------------------------------------------

Total Number of Routes : 2

10.1.1.0/24         (i) area: 0.0.0.0
     directly attached to interface 1/1/1, cost 1 distance 110
20.1.1.0/24         (I)
     via 10.1.1.2 interface 1/1/1, cost 2 distance 110
```

Showing OSPF routing table information for a specific subnet:

```
switch# show ip ospf routes 10.1.1.0/2
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1 VRF default, Routing Table for prefixes 10.1.1.0/24
---------------------------------------------------------------------

Total Number of Routes : 1

10.1.1.0/24         (i) area: 0.0.0.0
     directly attached to interface 1/1/1, cost 1 distance 110
```

Showing OSPF routing table information for a specific protected subnet on the default VRF:

```
switch# show ip ospf routes 211.0.0.0/24 protected vrf default
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1, VRF default, Protected Routes for prefixes 211.0.0.0/24
------------------------------------------------------------------------------

211.0.0.0/24       (i) area: 0.0.0.0
        via 110.110.110.1, interface vlan20, cost 122
        bkp 130.130.130.1, interface vlan22, cost 132


Total Number of Routes : 1
```

Showing OSPF routing table information for a specific protected subnet on all VRFs

```
switch# show ip ospf routes 211.0.0.0/24 protected all-vrfs
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1, VRF default, Protected Routes for prefixes 211.0.0.0/24
------------------------------------------------------------------------------

211.0.0.0/24       (i) area: 0.0.0.0
        via 110.110.110.1, interface vlan20, cost 122
        bkp 130.130.130.1, interface vlan22, cost 132


Total Number of Routes : 1
```

Showing OSPF routing table information for a specific protected subnet:

```
switch# show ip ospf routes 211.0.0.0/24 protected
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1, VRF default, Protected Routes for prefixes 211.0.0.0/24
------------------------------------------------------------------------------

211.0.0.0/24       (i) area: 0.0.0.0
        via 110.110.110.1, interface vlan20, cost 122
        bkp 130.130.130.1, interface vlan22, cost 132

Total Number of Routes : 1
```

Showing OSPF routing table information for all protected routes:

```
switch# show ip ospf routes 211.0.0.0/24 protected
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1, VRF default, Protected Routes
------------------------------------------------------------------------------

120.120.120.0/24       (i) area: 0.0.0.0
        via 110.110.110.1, interface vlan20, cost 120
        bkp 130.130.130.1, interface vlan22, cost 230
```

```
140.140.140.0/24       (i) area: 0.0.0.0
        via 110.110.110.1, interface vlan20, cost 130
        bkp 130.130.130.1, interface vlan22, cost 220

211.0.0.0/24       (i) area: 0.0.0.0
        via 110.110.110.1, interface vlan20, cost 122
        bkp 130.130.130.1, interface vlan22, cost 132

Total Number of Routes : 3

OSPFs Process ID 1, VRF red, Protected Routes
---------------------------------------------

Total Number of Protected Routes: 0
```

Showing OSPF routing table information for protected routes on the default VRF:

```
switch# show ip ospf routes protected vrf default
Codes: i - Intra-area route, I - Inter-area route
       E1 - External type-1, E2 - External type-2

OSPF Process ID 1, VRF default, Protected Routes
------------------------------------------------------------------------------

140.140.140.0/24       (i) area: 0.0.0.0
        via 130.130.130.1, interface vlan22, cost 130
        bkp 110.110.110.1, interface vlan20, cost 220

211.0.0.0/24      (i) area: 0.0.0.0
        via 110.110.110.1, interface vlan20, cost 122
        bkp 130.130.130.1, interface vlan22, cost 132

120.120.120.0/24       (i) area: 0.0.0.0
        via 110.110.110.1, interface vlan20, cost 120
        bkp 130.130.130.1, interface vlan22, cost 230

Total Number of Routes : 3
```

Command History

Release	Modification
10.15	Introduces the **protected** parameter.
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A	Operator (>) or Manager (#)	

Platforms	Command context	Authority
6400 8100 8325 8325H 8325P 8360 9300 9300S 10000		

show running-config bgp

```
show running-config bgp [vsx-peer]
```

Description

This command shows all configured BGP commands.

Parameter	Description
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Examples

EVPN only applies to HPE Aruba Networking 8100, 8325, 8360, 9300/9300S, 6300, 6400 Switch series.

```
switch# show running-config bgp
router bgp 65534.65535
    bgp asnotation dotted
    network 2.2.2.0/24
    neighbor 1.1.1.2 remote-as 65533.65535
    address-family ipv4 unicast
      neighbor 1.1.1.2 activate
      neighbor 1.1.1.2 route-map A out
      vrf v1
    address-family l2vpn evpn
      neighbor 1.1.1.2 activate
      neighbor 1.1.1.2 send-community extended
    exit-address-family
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

timers bgp

```
timers bgp <KEEPALIVE> <HOLDTIME>
no timers bgp <KEEPALIVE> <HOLDTIME>
```

Description

The command adjusts BGP network timers.

The **no** form of this command resets the BGP timers to defaults of 60 seconds for the keepalive timer and 180 seconds for the holdtime timer.

Parameter	Description
<KEEPALIVE>	Sets the value for keepalive timer. Default: 60 seconds. Range: 0-65535.
<HOLDTIME>	Sets the value for holdtime timer. Default: 180 seconds. Range: 0-65535.

Usage

- The keepalive timer is the number of seconds a BGP peer waits for a keep-alive message from a BGP peer before deciding the connection is down.

The holdtime timer is the number of seconds a BGP peer waits after not receiving a keepalive, update, or notification message before declaring that a connection with BGP peer is down.

- When a session is started, BGP negotiates holdtime with the neighbor, and selects the smaller value. The keepalive timer is then set based on the negotiated holdtime and the configured keepalive time.

Examples

```
switch(config-bgp)# timers bgp 100 150
switch(config-bgp)# no timers bgp
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-bgp`	Administrators or local user group members with execution rights for this command.

vrf

```
vrf <VRF-NAME>
no vrf <VRF-NAME>
```

Description

Creates a VRF instance named <*VRF-NAME*> and then enters its context. Use **default** for <*VRF-NAME*> to enter the default VRF configure context.

Except for the default VRF, the **no** form of the command deletes the named VRF instance and any IP configuration for interfaces or SVI linked to default VRF. The default VRF cannot be deleted and a warning is given if attempted. To erase the Route-Distinguisher and Route-Targets, enter the default VRF context and delete them manually one by one.

Only default and mangement VRFs are supported on the following 6300 switches: S3L75A, S3L76A, S3L77A.

Parameter	Description
<*VRF-NAME*>	Specifies the VRF name. Range: Up to 32 alphanumeric characters. The **mgmt** VRF cannot be used.

Examples

Creating the VRF named **cust_A** and then entering its context:

```
switch(config)# vrf cust_A
```

Entering the **default** VRF context:

```
switch(config)# vrf default
```

Deleting the VRF named **test**:

```
switch(config)# no vrf test
```

Command History

Release	Modification
10.09	Added default VRF information.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

IP Fast Re-Route (IP-FRR)

Overview

IP Fast Re-Route (IP-FRR) and BGP Prefix Independent Convergence (BGP-PIC) are both IP resiliency technologies to provide fast traffic recovery upon link or nexthop failure.

The goal is to reduce packet loss that happens while routers converge after a topology change due to a failure, as stated in RFC 5286 (Basic Specification for IP Fast Reroute: Loop-Free Alternates). To achieve fast traffic recovery, Routes are assigned primary nexthops and one backup nexthop. Routes that are assigned primary and backup nexthops are referred to as Protected Routes. Routes are programmed with the primary nexthops for traffic forwarding in the data plane. If the primary nexthops become unreachable, the Data Plane is reprogrammed to forward traffic to the backup nexthops. To achieve fast recovery, backup paths are precomputed so they can be activated as soon as the failure is detected. Routes that are not assigned with any alternative nexthops are referred to as Unprotected Routes (also referred to as native single and ECMP Routes).

Several technologies exist to determine when a nexthop become unreachable. Because of the requirement to quickly detect link failure between routers and switches, Bidirectional Forwarding Detection (BFD, RFC 5880) is used to detect when the primary nexthops are no longer reachable.

The OSPF routing protocol uses IP-FRR (RFC 5286) to achieve fast recovery. The BGP routing protocol uses BGP-PIC to achieve fast recovery.

BGP-PIC is a IETF draft document (draft-ietf-rtgwg-bgp-pic-19) and the feature is subject to change in future releases.

IP-FRR and BPG-PIC have separate configuration and show commands which are 58 described here.

Prerequisites

- At this time IP-FRR and BGP-PIC are supported on 9300 (Carmel) and 9300S (Paw-Patrol).
- Bidirectional Forwarding Detection (BFD) must be configured on the primary nexthops.
- OSPF point-to-point must be enabled on OSPF nexthops for IP FRR to compute protected routes.

Limitations

- Only egress interfaces of ROP, LAG, SVI and sub-interfaces are supported.
- Only IPv4 unicast is supported, IPv6 is not supported.
- IPv4 multicast is not supported.
- Protected Routes can have only one primary and one backup nexthop.
- OSPF IP FRR assumes both the primary and the backup nexthops are in the same routing area.

Fast Traffic Recovery commands

fast-reroute

[no] fast-reroute

Description

This command enables IP Fast Re-Route (IP-FRR) and is available in the router OSFP context.

Parameter	Description
no	*Optional*. Disables IP-FRR.

Examples

The following example shows how to enable IP fast-reroute for this OSPF instance:

```
switch(config)# router ospf 1
switch(config-ospf-1)# fast-reroute
switch(config-ospf-1)#
```

The following example shows how to disable IP-FRR:

```
switch(config)# router ospf 1
switch(config-ospf-1)# no fast-reroute
switch(config-ospf-1)#
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300 9300S	`config-ospf-1`	Administrators or local user group members with execution rights for this command.

Fast Traffic Recovery commands

neighbor

```
neighbor <ipv4-address | ipv6-address | peer-group-name> fall-over bfd
```

Description

The BGP Support for BFD enables BGP to register with Bidirectional Forwarding Detection (BFD) to receive fast peering session deactivation messages from BFD.

You can either configure BFD support for BGP per neighbor or peer-group. Multi-hop BFD is not supported for BGP.

Parameter	Description
ipv4-address	*Required*. The neighbor IP address. Syntax: A.B.C.D
ipv6-address	*Required*. The neighbor IPv6 address. Syntax:X:X::X:X
peer-group-name	*Required*. Configure Peer-Group. Syntax: PGNAME
fall-over	*Required*. Configure BGP session fast deactivation.

Example

The following example shows how configure BGP session fast deactivation and enable BFD for the BGP neighbor:

```
switch(config)# router bgp 1
switch(config-bgp)# neighbor 1.1.1.1 fall-over bfd
switch(config-bgp)# neighbor 1000::1 fall-over bfd
switch(config-bgp)# neighbor PG fall-over bfd
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300 9300S	`config-ospf-1`	Administrators or local user group members with execution rights for this command.

show capacities fast-reroute

show capacities fast-reroute

Description

Displays the maximum number of Protected Nexthops configurable in a system.

Example

The following example shows the maximum Protected Nexthop configurable in a system:

```
switch# show capacities fast-reroute
System Capacities: Filter Fast-Reroute
Capacities Name                                                        Value
-----------------------------------------------------------------------------
Maximum number of Protected Nexthops configurable in a system          256
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300 9300S	Manager (#)	Administrators or local user group members with execution rights for this command.

show capacities status fast-reroute

show capacities-status fast-reroute

Description

Displays the total number of Protected Nexthops currently configured on the device.

Example

The following example shows the currently configured and the maximum Protected Nexthop in a system:

```
switch# show capacities-status fast-reroute

System Capacities Status: Filter Fast-Reroute
Capacities Name                                                  Value Maximum
-----------------------------------------------------------------------------
Maximum number of Protected Nexthops currently configured          2     256
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300 9300S	Manager (#)	Administrators or local user group members with execution rights for this command.

Chapter 8

Route Policies and Route Maps

Overview

Route policies control routing paths by filtering and modifying routing information using route maps. Routing policies can filter advertised, received, and redistributed routes. They can also modify attributes for specific routes.

Configuring a route policy is two-step process:

1. Configure filters based on route attributes, such as destination address and the advertising router address.
2. Create a route map and apply filters to it.

Route maps

Route maps can be used for route redistribution. Route map entries consist of a list of match and set criteria. Match criteria specify match conditions for incoming routes. Set criteria specify the action taken if match criteria are met.

Users can configure multiple entries in the same route map. These entries contain the same route map name and are differentiated by a sequence number. A sequence number will automatically be assigned if user does not provide one.

Create a route map with one or more route map entries arranged by the sequence number under a unique route map name. The route map entry has the following parameters:

- Sequence number
- Permission (permit or deny)
- Match criteria
- Set changes

By default, a route map processes routes starting from the lowest sequence number to the highest. Other filtering mechanisms include Prefix List, AS Path List, and Community List.

Match criteria

You can use a variety of criteria to match a route in a route map. Some criteria, such as BGP community lists and AS-Path Lists, are applicable only to a specific routing protocol. Other criteria, such as the Prefix lists, can be used for any route.

When a routing protocol processes a route through a route map, it compares the route to each of the configured match statements. If the route or packet matches the configured criteria, the routing protocol processes it based on the permit or deny configuration for that match entry in the route map.

The match categories and parameters are as follows:

- **BGP parameters**: Match based on AS-path Lists, Community Lists, or Origin code.
- **Prefix lists**: Match based on an address or range of addresses.
- **Other parameters**: Match based on an IP next-hop address or metric.

Set changes

Once a route matches an entry in a route map, the route can be changed based on one or more configured set statements.

Set changes are as follows:

- **BGP parameters**: Change the AS-path, community, local preference, origin, or weight attributes.
- **Metrics**: Change the route metric.
- **Other parameters**: Change the forwarding address or the IP next-hop address.

IP prefix lists

Use prefix lists to permit or deny an address or range of addresses. Filtering by a prefix list involves matching the prefixes of routes with the prefixes listed in the prefix list. An implicit deny is assumed if a prefix does not match any entries in a prefix list. Prefix lists can be configured for both IPv4 and IPv6 addresses.

Users can configure multiple entries in a prefix list and permit or deny the prefixes that match the entry. Each entry has an associated sequence number.

AS-path lists for BGP

Use AS-path lists to filter inbound or outbound BGP route updates. If the route update contains an AS-path attribute that matches an entry in the AS-path list, the router processes the route based on the configured permit or deny condition.

Configure multiple AS-path entries in an AS-path list by using the same AS-path list name. The router processes the first matching entry. You can also configure an AS-path list to compare the AS numbers against a regular expression.

Community lists for BGP

Filter BGP route updates based on the BGP community attribute by using community lists in a route map. Match the community attribute based on a community list and set the community attribute using a route map.

A community list contains one or more community attributes. If a user configures more than one community attribute in the same community list entry, the BGP route must match all listed community attributes to be considered a match.

Configure multiple community attributes as individual entries in the community list by using the same community list name. The router processes the first community attribute that matches the BGP route, using the permit or deny configuration for that entry.

Configure community attributes in the community list in one of the following formats:

- A named community attribute, such as Internet or no-export.
- In aa:nn format.

Route flap dampening

To maintain scalability of a routed internet, it is necessary to reduce the amount of change in routing state propagated by BGP in order to limit processing requirements. The primary contributors of processing load resulting from BGP updates are the BGP decision process and adding and removing forwarding entries.

A route flap occurs when the state of a route changes from up to down or down to up. When a route state changes, the route tables of the devices that support the route also change. Route flap dampening prevents BGP from selecting unstable routes. If an EBGP peer goes down after you configure this feature, routes coming from the peer are dampened but not deleted.

This feature applies to EBGP routes but not IBGP routes.

Route redistribution and route maps

Use route maps to control the redistribution of routes between routing domains. Route maps match on the attributes of the routes to redistribute routes that pass the match criteria. The route map can also modify the route attributes using set changes.

The router matches redistributed routes against each route map entry.

- If there are multiple match statements, the route must pass all match criteria.
- If a route passes the match criteria defined in a route map entry, actions defined in the entry are executed.
- If the route does not match the criteria, the router compares the route against subsequent route map entries.

Route processing continues until a match is made or the route is processed by all entries in the route map with no match. If the router processes the route against all entries in a route map with no match, the router accepts the route (inbound route maps) or forwards it (outbound route maps).

Route policy and route map commands

This section describes general and filtering commands, as well as match, set and show commands for configuring route policies and route maps.

General or filtering commands

ip aspath-list

```
ip aspath-list <ASPATH-LIST-NAME> [seq <SEQ>] {permit | deny} <REGEXP>
no ip aspath-list <ASPATH-LIST-NAME> [seq <SEQ>] {permit | deny} <REGEXP>
```

Description

Configures an AS Path list to match a specific AS path. AS Path lists are named lists of regular expression rules. They are used to match AS Path attributes in the routes for inclusion in or exclusion from route

policies. The sequence number is optional and is autogenerated whenever it is not explicitly mentioned. All AS Path list rules with the same name are grouped together.

The **no** form of this command removes the AS Path list configuration.

Parameter	Description
<ASPATH-LIST-NAME>	Specifies the name of the AS Path list.
`seq` *<SEQ>*	Specifies the order of reference of the regular expression rules.
`{permit | deny}`	Specifies whether the route is available for further processing when there is a match.
<REGEXP>	Specifies the regular expression to match the AS Path. Standard regular expression wildcards are supported. The _ character can be used to match the AS Path boundary.

Examples

Configuring an AS Path list with sequence numbering:

```
switch(config)# ip aspath-list ASLst seq 5 permit _4*
```

Configuring a prefix list without sequence numbering:

```
switch(config)# ip aspath-list ASLst permit _4*
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

ip community-list

`ip community-list standard` *<COMM-LIST-NAME>* *<DESCRIPTION>* `[seq` *<SEQ>*`] {permit | deny}` *<COMMUNITY-NUMBER>*

```
no ip community-list standard <COMM-LIST-NAME> <DESCRIPTION> [seq <SEQ>] {permit | deny}
<COMMUNITY-NUMBER>
```

Description

Configures a community list to match a specific community number attribute. Community-list is a named list of regular expressions. They are used to match the community number attributes in the routes for inclusion in, or exclusion from route policies. The sequence number is optional and is autogenerated whenever it is not explicitly mentioned. All community-list rules with the same name are grouped.

The **no** form of this command removes the community list configuration.

Parameter	Description
<COMM-LIST-NAME>	Specifies the name of the community list that matches community number of a route.
<DESCRIPTION>	Specifies the description of the community list. Maximum character limit is 80.
`seq` *<SEQ>*	Specifies the order of reference of the regular expression rules.
`{permit | deny}`	Specifies whether the route is available for further processing when there is a match.
<COMMUNITY-NUMBER>	Specifies the community number. The community number must be in AA: NN format or from the list of well-known community.

Examples

Configuring a community list with sequence numbering:

```
switch(config)# ip community-list standard CommLst seq 5 permit 101:41
```

Configuring a community list without sequence numbering:

```
switch(config)# ip community-list standard CommLst no-export permit
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400	`config`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300 9300S 10000		

ip prefix-list

```
ip prefix-list <PREFIX-LIST-NAME> [seq <SEQ>] <IP-PREFIX/MASK> [ge <0-32>] [le <0-32>]
no ip prefix-list <PREFIX-LIST-NAME> [seq <SEQ>] <IP-PREFIX/MASK> [ge <0-32>] [le <0-32>]
```

Description

Configures a prefix list to match a set of prefixes. Prefix lists are named lists of route prefixes. They are used to match routes for inclusion in or exclusion from route policies. The sequence number determines the order of matching. The matches are performed starting from the lowest sequence number to the highest sequence number until there is a match. The sequence number is however optional and is autogenerated whenever it is not explicitly mentioned. All prefixes with the same prefix list name are grouped.

The autogenerated sequence number is derived by adding 10 to the highest sequence number available. This technique makes it possible to insert new prefix list sequence number in between.

The **ge** and **le** parameters are used to combine prefixes with a range of network mask. For example, **172.131.0.0/16 ge 16 le 24** will match all prefixes within the 172.131.0.0/16 network that have a mask greater than or equal to 16 bits and less than or equal to 24 bits in length. For instance, 172.131.1.0/18 would match, because its length is between 16 and 24 but 172.0.0.0/8 or 172.131.1.128/25 would not match.

The **no** form of this command removes the prefix list configuration. Prefix-list commands which generate sequence numbers must explicity use sequence numbers in the **no** form.

Parameter	Description
<PREFIX-LIST-NAME>	Specifies the name of the prefix list.
seq *<SEQ>*	Specifies the order of reference of the prefix rules.
{permit | deny}	Specifies whether the route is available for further processing when there is a match.
IP-PREFIX/MASK>	Specifies the IP prefix or mask.
ge *<0-32>*	Specifies the minimum prefix length to be matched.
le *<0-32>*	Specifies the maximum prefix length to be matched.

Examples

Configuring a prefix list with sequence numbering:

```
switch(config)# ip prefix-list PFXLST seq 5 permit 4.0.0.0/8 ge 9 le 12
```

Configuring a prefix list without sequence numbering:

```
switch(config)# ip prefix-list PSXLST permit 5.0.0.0/8
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config	Administrators or local user group members with execution rights for this command.

ipv6 prefix-list

```
ipv6 prefix-list <PREFIX-LIST-NAME> [seq <SEQ>] <IPV6-PREFIX/MASK> [ge <0-128>] [le <0-128>]
no ipv6 prefix-list <PREFIX-LIST-NAME> [seq <SEQ>] {ip | ipv6} <IPV6-PREFIX/MASK> [ge <0-128>] [le <0-128>]
```

Description

Configures a prefix list to match a set of prefixes. Prefix lists are named lists of route prefixes. They are used to match routes for inclusion in or exclusion from route policies. The sequence number determines the order of matching. The matches are performed starting from the lowest sequence number to the highest sequence number until there is a match. The sequence number is however optional and is autogenerated whenever it is not explicitly mentioned. All prefixes with the same prefix list name are grouped.

The autogenerated sequence number is derived by adding 10 to the highest sequence number available. This technique makes it possible to insert new prefix list sequence number in between.

The **ge** and **le** parameters are used to combine prefixes with a range of network mask. For example, **2000::/64 ge 65 le 70** will match all prefixes within the 2000::/64 network that have a mask greater than or equal to 65 bits and less than or equal to 70 bits in length.

The **no** form of this command removes the prefix list configuration. Prefix-list commands which generate sequence numbers must explicity use sequence numbers in the **no** form.

Parameter	Description
<PREFIX-LIST-NAME>	Specifies the name of the prefix list.
`seq` *<SEQ>*	Specifies the order of reference of the prefix rules.
`{permit | deny}`	Specifies whether the route is available for further processing when there is a match.
IP-PREFIX/MASK>	Specifies the IP prefix or mask.
`ge` *<0-128>*	Specifies the minimum prefix length to be matched.
`le` *<0-128>*	Specifies the maximum prefix length to be matched.

Examples

Configuring a prefix list with sequence numbering:

```
switch(config)# ipv6 prefix-list PFXLST seq 10 permit 2000::64 ge 65 le 70
```

Configuring a prefix list without sequence numbering:

```
switch(config)# ipv6 prefix-list PSXLST permit 2000::1/128
```

Removing the configuring of a prefix list:

```
switch(config)# no ipv6 prefix-list P2 seq 10 permit any
```

Removing the configuring of a prefix list:

```
switch(config)# no ipv6 prefix-list P1 seq 10
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H	`config`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325P 8360 9300 9300S 10000		

route-map

```
route-map <NAME> {permit | deny} seq <NUMBER>
no route-map <NAME> {permit | deny} seq <NUMBER>
```

Description

Configures a route map entry with the given name and action by taking the CLI in the route map context. All route map entries with the same name belong to the same route map. The route map entry rules are processed in order by sequence number, until a match is found.

The **no** form of this command removes the route map entry configuration.

Parameter	Description
<NAME>	Specifies the name of the route map. Required.
`{permit | deny}`	Specifies whether the route is available for further processing when there is a match. Required.
<NUMBER>	Specifies the sequence number of the entry. Required.

Examples

Configuring a route map entry:

```
switch(config)# route-map GlobalMap permit seq 10
```

Removing a route map entry configuration:

```
switch(config)# no route-map GlobalMap permit seq 10
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A	`config`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
6400 8100 8325 8325H 8325P 8360 9300 9300S 10000		

continue

```
continue <SEQUENCE NUMBER>
no continue
```

Description

Allows you to execute additional entries in a route map. The sequence number specifies the route-map entry's sequence number that will be executed next if the existing entry's match clause is successful.

If a successful match occurs and continue command exists, the route map saves the set value first and then jumps to the specified route map entry.

Set clauses are saved during the match clause evaluation and are executed only after the route map evaluation is completed. The set clauses are executed in the order in which they were configured.

Set clauses can be accumulative or additive as **set as-path prepend** or it can be absolute as **set metric**. For set commands that configures an accumulative value, subsequent values are added in order in which they were configured. For set commands that configures an absolute value, The values from the last instance will be applied.

The **no** form of this command removes the route map continue configuration.

If the specified route-map sequence entry does not exist, route-map processing will be terminated at the current sequence number if its clause is matched. The continue sequence number must be higher than the current route map sequence number for this command to take effect.

Parameter	Description
<SEQUENCE NUMBER>	Specifies the value of the route map entry to be executed next after a successful match clause.

Examples

Configuring a route map to continue to execute an additional entry:

```
switch(config)# route-map GlobalMap permit 10
switch(config-route-map-GlobalMap-10)# continue 40
```

Removing a route map continue configuration:

```
switch(config)# route-map GlobalMap permit seq 10
switch(config-route-map-GlobalMap-10)# no continue 40
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

Match commands

match aspath-list

```
match aspath-list <ASPATH-LIST-NAME>
no match aspath-list <ASPATH-LIST-NAME>
```

Description

Matches the AS path attribute of the route with one or more regular expressions in the AS path list.

The **no** form of this command restores the default behavior of not matching the AS path attribute of the route.

Parameter	Description
<ASPATH-LIST-NAME>	Specifies the name of the AS path list to match the AS path attribute of the route.

Example

Configuring a match clause in the route map to match the AS path list:

```
switch(config)# ip aspath-list ASLst permit 1001
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# match aspath-list ASLst
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-route-map`	Administrators or local user group members with execution rights for this command.

match community-list

```
match community-list <COMMUNITY-LIST-NAME> [exact-match]
no match community-list <COMMUNITY-LIST-NAME> [exact-match]
```

Description

Matches the community number attribute of the route with one, or more regular expressions in the community-list.

The **no** form of this command restores the default behavior of not matching the community number attribute of the route.

Parameter	Description
<COMMUNITY-LIST-NAME>	Specifies the name of the community-list to match the community number attribute of the route.
`[exact-match]`	Indicates that the community number attribute must match exactly with the expressions in the community-list. However, the order of the communities in the community-list is of no significance.

Example

Configuring a match clause in the route map to match the community list:

```
switch(config)# ip community-list standard CommLst 101:41 permit 12:201
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# match community-list CommLst
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-route-map	Administrators or local user group members with execution rights for this command.

match interface

```
match interface <INTERFACE-NAME>
no match interface <INTERFACE-NAME>
```

Description

Matches the outgoing interface value of the route with the value configured in the match clause. This command is applicable to static and connected routes which will be redistributed to the BGP protocol.

The **no** form of this command restores the default behavior of not matching the outgoing interface value of the route.

Parameter	Description
<INTERFACE-NAME>	Specifies the value to be matched with the outgoing interface of the route entry.

Example

Configuring a match clause in the route map to match the outgoing interface of the route:

```
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# match interface  1/1/1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400	config-route-map	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300 9300S 10000		

match ip address prefix-list

```
match ip address prefix-list <PREFIX-LIST-NAME>
no match ip address prefix-list <PREFIX-LIST-NAME>
```

Description

Matches the destination IP address prefix of the routes with one or more addresses in the prefix list.

The **no** form of this command restores the default behavior of not matching the destination IP address prefix of the routes to their default value. This command is applicable to OSPF, static, and connected routes which will be redistributed to the BGP protocol.

Parameter	Description
<PREFIX-LIST-NAME>	Specifies the name of the prefix list to be matched with the network address of the route.

Example

Configuring a prefix list and a match clause in route map to match the prefix list:

```
switch(config)# ip prefix-list PfxLst permit 4.0.0.0/8
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# match ip address prefix-list PfxLst
```

When the IP prefix list with prefix and mask-length of *0.0.0.0/0* is used, the route matches default-route *0.0.0.0/0* as well as any other route. This behavior would be changed to match only the default-route in the next release.

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A,	`config-route-map`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000		

match ip next-hop

```
match {ip | ipv6} next-hop {<ADDRESS> | prefix-list <PREFIX-LIST-NAME>}
no match {ip | ipv6} next-hop [<ADDRESS> | prefix-list <PREFIX-LIST-NAME>]
```

Description

Matches the next-hop address of the route with the configured address in the match clause. This command is applicable to static routes which will be redistributed to BGP protocol.

The **no** form of this command restores the default behavior of not matching the next-hop address of the route.

Parameter	Description
<ADDRESS>	Specifies the IPv4 address to match with the next-hop address of the route.
`prefix-list` *<PREFIX-LIST-NAME>*	Specifies the name of the IP prefix list to be matched with the next-hop address of the route.

Example

Configuring a match clause in the route map to match the next-hop address of the route:

```
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# match ip next-hop 1.1.1.2
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A	`config-route-map`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
6400 8100 8325 8325H 8325P 8360 9300 9300S 10000		

match ip route-source

```
match ip route-source prefix-list <PREFIX-LIST-NAME>
no match ip route-source prefix-list <PREFIX-LIST-NAME>
```

Description

Matches the IP address of the source of the route using IP prefix lists.

The **no** form of this command restores the default behavior of not matching the IP address of the route.

Parameter	Description
<PREFIX-LIST-NAME>	Specifies the name of the prefix list to match the IP address of the source of the route.

Example

Configuring a match clause in the route map to match the source of the route:

```
switch(config)# ip prefix-list RouterLst 4.4.4.4/32
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# match ip route-source prefix-list RouterLst
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360	config-route-map	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
9300 9300S 10000		

match local-preference

```
match local-preference <VALUE>
no match local-preference <VALUE>
```

Description

Matches the local preference value of the route with the value configured in the match clause.

The **no** form of this command restores the default behavior of not matching the local preference value of the route.

Parameter	Description
<VALUE>	Specifies the value to be matched with the route entry local preference in the range of 1 to 4294967295.

Example

Configuring a match clause in the route map to match the local preference of the route:

```
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# match local-preference 100
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-route-map	Administrators or local user group members with execution rights for this command.

match metric

```
match metric <VALUE>
```

```
no match metric <VALUE>
```

Description

Matches the MED value of the route with the value configured in the match clause.

The **no** form of this command restores the default behavior of not matching the MED value of the route.

Parameter	Description
<VALUE>	Specifies the value to be matched with the route entry MED.

Example

Configuring a match clause in the route map to match the metric of the route:

```
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# match metric 10
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-route-map	Administrators or local user group members with execution rights for this command.

match origin

```
match origin {igp | egp | incomplete}
no match origin [igp | egp | incomplete]
```

Description

Matches the route origin attribute of the route with route configured in the match clause.

The **no** form of this command restores the default behavior of not matching the route origin attribute of the route.

Parameter	Description
`{igp | egp | incomplete}`	Specifies if the route origin attribute is matched with a match clause which originated as IGP, EGP, or has unknown origin. The unknown origin is typically redistributed from another routing protocol.

Example

Configuring a match clause in the route map to match the origin:

```
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# match origin igp
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-route-map`	Administrators or local user group members with execution rights for this command.

match route-type

```
match route-type external {type-1 | type-2}
no match route-type external [type-1 | type-2]
```

Description

Matches the metric-type value of the OSPF external route with the value configured in the match clause.

The **no** form of this command restores the default behavior of not matching the metric-type value of the OSPF external route.

Parameter	Description
`{type-1 | type-2}`	Specifies the *type-1* or *type-2* OSPF value to be matched with the external route.

Example

Configuring a match clause in the route map to match the metric-type value of the OSPF external route:

```
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# match route-type external type-1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-route-map`	Administrators or local user group members with execution rights for this command.

match source-protocol

```
match source-protocol {bgp | connected | ospf | static}
no match source-protocol [bgp | connected | ospf | static]
```

Description

Matches the source routing protocol value of the route with the value configured in the match clause.

The **no** form of this command restores the default behavior of not matching the source routing protocol value of the route.

Parameter	Description
`{bgp | connected | ospf | static}`	Specifies the *bgp*, *connected*, *ospf*, or *static* value to be matched with the route entry source protocol.

Example

Configuring a match clause in the route map to match the source protocol route:

```
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# match source-protocol ospf
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-route-map`	Administrators or local user group members with execution rights for this command.

match tag

```
match tag <value>
no match tag <value>
```

Description

Matches the tag value of the route with the one configured in the match clause. Applies to static routes that will be redistributed to BGP, OSPFv3 and OSPFv3 protocols.

The **no** form of this command removes the tag value of the route.

Parameter	Description
`value`	Numeric value to match with the route tag. Required.

Example

Configuring a match clause in route-map to match the tag value of the route:

```
switch(config)# route-map GlobalMap permit 11
switch(config-route-map)# match tag 20
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-route-map`	Administrators or local user group members with execution rights for this command.

match vni

```
match vni <value>
no match vni <value>
```

Description

Matches the VNI value of the route with the one configured in the `match` clause. Applies to matching L2VNIs or L3VNIs. Use the `continue` clause if both L2VNI and L3VNI are to be matched. Route maps with the `match vni` clause can be used with L2VPN EVPN neighbors only.

The **no** form of this command removes the match for the VNI value of the route.

Parameter		Description
	<value>	Numeric value to match with the route tag. Required. Range: 1 to 16777214

Example

Configuring a match clause in route map to match the VNI value of the route:

```
switch(config)# route-map GlobalMap permit 11
switch(config-route-map)# match vni 10000
```

Configuring a match clause in route map to match both L2VNI and L3VNI in a single route map:

```
switch(config)# route-map GlobalMap permit 11
switch(config-route-map)# match vni 10000
switch(config-route-map)# continue 12
switch(config)# route-map GlobalMap permit 12
switch(config-route-map)# match vni 10
```

The following example is different from the one above, as it configures a match clause in route map to match any one of the two VNIs:

```
switch(config)# route-map GlobalMap permit 10
switch(config-route-map)# match vni 10000
switch(config)# route-map GlobalMap permit 20
```

```
switch(config-route-map)# match vni 10
switch(config)# route-map GlobalMap deny 30
```

Command History

Release	Modification
10.09	Command introduced

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8325 8325H 8325P 8360 9300 9300S 10000	`config-route-map`	Administrators or local user group members with execution rights for this command.

Set commands

set as-path exclude

```
set as-path exclude <AS>
no set as-path exclude <AS>
```

Description

Removes all occurrences of the configured AS Path from the AS Path attribute of the route.

The **no** form of this command restores the default behavior of not modifying the AS Path attribute list.

Parameter	Description
<AS>	Specifies the AS number to be removed from the AS Path attribute of the route.

Example

Configuring a set clause in the route map to remove the AS from the AS Path attribute of the route:

```
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# set as-path exclude 1001
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-route-map`	Administrators or local user group members with execution rights for this command.

set as-path prepend

```
set as-path prepend <AS> <AS>...
no set as-path prepend <AS> <AS>...
```

Description

Pretends the list of the configured AS numbers to the AS Path attribute of the routes. To ensure that the AS path conforms to standards, the local AS is prepended after this command is executed.

The **no** form of this command restores the default behavior of not modifying the AS Path attribute list.

Parameter	Description
<AS> <AS>...	Specifies the AS numbers to be prepended from the AS Path attribute of the route.

Example

Configuring a set clause in the route map to prepend the AS from the AS Path attribute of the route:

```
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# set as-path prepend 1
switch(config-route-map-GlobalMap-11)# no set as-path prepend 102
```

The **no** form of the command deletes the entire list of AS-Path prepend configuration regardless of the parameter list. In this example, the **no** form command would result in deletion of all the three AS numbers that were earlier configured.

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-route-map`	Administrators or local user group members with execution rights for this command.

set community

```
set community {<AA:NN> | internet | no-export | no-advertise | local-as} [additive | delete]
no set community [<AA:NN> | internet | no-export | no-advertise | local-as] [additive | delete]
```

Description

Modifies the community number attribute of the route with the value configured in the set clause. This command is applicable to OSPF, static, and connected routes which will be redistributed to the BGP protocol.

The **no** form of this command restores the default behavior of not modifying the community number attribute of the route.

Parameter	Description
`{<AA:NN> | internet | no-export | no-advertise | local-as}`	Selects the value to be set as the community number attribute of the route in the **AA:NN** format (quotation marks required when multiple communities are listed, for example: **set community "65001:100 65001:200"**) or as a known community name **internet**, **no-export**, **no-advertise**, and **local-as**.
`[additive]`	Specifies that the specified community number is added to the existing community number attribute of the route.
`[delete]`	Specifies that the specified community number is removed from the existing community number attribute of the route.

Example

Configuring a set clause in the route map to modify the community number attribute of the route:

```
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# no set community 11:101
switch(config-route-map-GlobalMap-11)# set community no-advertise
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-route-map	Administrators or local user group members with execution rights for this command.

set dampening

```
set dampening {half-life <VALUE> reuse <VALUE> suppress <VALUE> max-suppress-time
<VALUE>}
no set dampening
```

Description

Sets parameters of route flap dampening feature.

Parameter	Description
half-life	Time to reduce the penalty to half.
reuse	Lower threshold of penalty.
suppress	Upper threshold of penalty.
max-suppress-time	Max time to keep route suppressed.

Example

```
switch(config-route-map-abc-20)# set dampening half-life 5 reuse 50 suppress 125
max-suppress-time 255

switch(config-route-map-abc-20)# no set dampening half-life 5 reuse 50 suppress
125 max-suppress-time 255
```

```
switch(config-route-map-abc-20)# set dampening

switch(config-route-map-abc-20)# no set dampening
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-route-map`	Administrators or local user group members with execution rights for this command.

set extcommunity

```
set extcommunity "[rt <VALUE> | evpn-rmac <MAC-ADDRESS>]"
no set extcommunity [rt | evpn-mac]
```

Description

Sets the extended community number attribute for a route matching the route map. Extended communities are supported only on Route Targets. This command is applicable to OSPF, static, and connected routes which will be redistributed to the BGP protocol.

The **no** form of this command restores the default behavior of not modifying the extended community number attribute of the route.

Parameter		Description
	<VALUE>	Sets the extended community number attribute. Specify the information in asn:nn format.
	<MAC-ADDRESS>	Specifies MAC address of the Router-MAC extended community. When configuring a route map and using the **set ip next-hop** command. this value can be set for the OUT direction only.

Usage

- Multiple community numbers can be configured within the double quotes.
- 2-byte and 4-byte ASN values are supported in the global administrator component of the extended community attribute.
- 4-byte ASN values must be within the range of 1-4294967295.
- 4-byte ASN values do not support dotted notation.
- Extended communities are only supported on route targets.
- The

Examples

Configuring a set clause in a route-map to modify the community number attribute of the route:

```
switch(config)# route-map abc permit seq 10
switch(config-route-map-abc-10)# set extcommunity rt "1:1 2:2"
```

Configuring a set clause in a route-map to modify the router mac:

```
switch(config)# route-map abc permit seq 1
switch(config-route-map-abc-1)# set extcommunity evpn-mac 00:01:01:90:90:01
```

Command History

Release	Modification
10.10	Command introduced

Command Information

Platforms	Command context	Authority
6400 8100 8325 8325H 8325P 8360	config-route-map-abc-10	Administrators or local user group members with execution rights for this command.

set ip nexthop

```
set {ip | ipv6} nexthop {global} <IP-ADDR>
no set {ip | ipv6} nexthop {global} <IP-ADDR>
```

Description

Sets the IP address of the next-hop of the route with the value configured in the set clause.

The **no** form of this command restores the default behavior of not modifying the IP address of the next-hop of the route.

When a next-hop is changed via a route-map using the **set ip next-hop** command, the router-mac of the route remains unchanged. The router-mac of the route is changed/retained based on the **next-hop-self** and **next-hop-unchanged** configurations *only* if the next-hop is not changed via a route-map.

Parameter	Description
<IP-ADDR>	Specifies the IPv4 address to be set as the next-hop address of the route.

Example

Configuring a set clause in the route map to modify the next-hop address of the route entry:

```
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# set ip nexthop 1.1.1.2
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-route-map`	Administrators or local user group members with execution rights for this command.

set ipv6 nexthop global

```
set ipv6 nexthop global <IP-ADDRESS>
no set ip nexthop global <IP-ADDRESS>
```

Description

Sets the IPv6 address of the nexthop of the routes with the IPv6 address configured in the set clause.

The **no** form of this command removes this configuration.

Parameter	Description
<IP-ADDRESS>	Specifies the IPv6 address of the nexthop router.

Examples

Configuring a set clause in route-map to modify the nexthop address of route entry:

```
switch(config)# route-map GlobalMap permit 11
switch(config-route-map)# set ipv6 nexthop global 1.1.1.2
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360	`config`	Administrators or local user group members with execution rights for this command.

set local-broadcast-group

```
set local-broadcast-group <group ID>
no set local-broadcast-group
```

Description

Implements split-horizon of extended VLANs on Border-VTEPs in multi-fabric VXLAN deployments.

Parameter	Description
`group ID`	Specifies the group ID . **NOTE:** The supported range for broadcast group ID is 1 to 4 for 8325 switches **NOTE:** The supported range for broadcast group ID is 1 to 128 for 8360 switches

Usage

The **set local-broadcast-group** command is available in the route map context and is applicable only to BGP neighbors that are configured with the EVPN address family. The local-broadcast-group should not be used for any other protocol and route distribution. The local-broadcast-group is supported only for route maps applied with in-bound traffic direction. If applied for route-maps with out-bound traffic direction, the configuration will have no effect.

The **local-broadcast-group** cannot be configured when the **dyn-vxlan-tunnel-bridging-mode** is configured. If both the local-broadcast-group and dyn-vxlan-tunnel-bridging-mode are configured, an error message will be displayed and the user should remove one of the configurations and then issue

the **clear bgp** command. Similarly, the dyn-vxlan-tunnel-bridging-mode cannot be configured even when a single route map with local-broadcast-group exists without being applied to a BGP neighbor.

Route maps with local-broadcast-group should be configured only on VTEPs. If a VTEP configuration is removed from Border Gateway Protocol, the local-broadcast-group applied on BGP neighbors with the EVPN adress family should also be removed.

When the local-broadcast-group value is modified, users should issue the **clear bgp** command for the value change to take effect. This is required irrespective of the whether the VXLAN interface is up or down.

Multiple route maps can share the same broadcast-group ID as per the requirements of the deployment. There is no limitation to restrict a broadcast-group ID to a particular route map. Use caution while configuring the broadcast group.

Configuration of static tunnels is not recommended on VTEPs where split horizon is disabled between the EVPN tunnels, either by using the ibgp-ebgp command or by using the route map- based broadcast-group configuration.

Example

Configuring a set clause in the route map to modify the broadcast group ID:

```
switch(config)# route-map GlobalMap permit 11
switch(config-route-map-GlobalMap-11)# set local-broadcast-group 2
```

Command History

Release	Modification
10.13	Command Introduced

Command Information

Platforms	Command context	Authority
8325 8325H 8325P 8360	config-route-map	Administrators or local user group members with execution rights for this command.

set local-preference

```
set local-preference <VALUE>
no set local-preference <VALUE>
```

Description

Modifies the local-preference attribute of the route entry with the value configured in the set clause.

The **no** form of this command restores the default behavior of not modifying the local-preference attribute of the route entry.

Parameter	Description
<VALUE>	Specifies the value to be set as the local-preference attribute of the route entry. Range: 0 to 4294967295.

Example

Configuring a set clause in the route map to modify the metric value of the route:

```
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# set local-preference 100
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-route-map`	Administrators or local user group members with execution rights for this command.

set metric

```
set metric <VALUE>
no set metric <VALUE>
```

Description

Modifies the metric value of the route with the value configured in the set clause. This command is applicable to OSPF, static, and connected routes which will be redistributed to the BGP protocol.

The **no** form of this command restores the default behavior of not modifying the metric value of the route.

Parameter	Description
<VALUE>	Specifies the value to be set as the metric value of the route. Range: 0 to 4294967295.

Example

Configuring a set clause in the route map to modify the metric value of the route:

```
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# set metric 10
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-route-map`	Administrators or local user group members with execution rights for this command.

set origin

```
set origin [igp | egp | incomplete]
no set origin [igp | egp | incomplete]
```

Description

Modifies the route origin attribute of the route update with the value configured in the set clause.

The **no** form of this command restores the default behavior of not modifying the route origin attribute of the route.

Parameter	Description
`{igp | egp | incomplete}`	Selects the route update originated to IGP, EGP, or incomplete. When incomplete is selected, the route update origin is set to unknown.

Example

Configuring a set clause in the route map to modify the origin attribute of the route:

```
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap-11)# set origin igp
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-route-map`	Administrators or local user group members with execution rights for this command.

set tag

```
set tag <value>
[no] set tag <value>
```

Description

Modifies the tag value of the route with the one configured in the set clause. Applicable to static routes that will be redistributed to ospfv2 and ospfv3 protocols.

The **no** form of this command removes the set clause tag value.

Parameter	Description
`value`	Numeric value to change the route entry tag. Range: 0-4294967295. Required.

Example

Configuring a set clause in route-map to modify the tag value of the route:

```
switch(config)# route-map GlobalMap permit 11
switch(config-route-map)# set tag 10
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-route-map`	Administrators or local user group members with execution rights for this command.

set weight

```
set weight <VALUE>
no set weight <VALUE>
```

Description

Modifies the weight attribute of the route with the value configured in the set clause.

The **no** form of this command restores the default behavior of not modifying the weight attribute of the route.

Parameter	Description
<VALUE>	Specifies the value to be set as the weight attribute of the route. Range: 0 to 65535.

Example

Configuring a set clause in the route map to modify the metric value of the route:

```
switch(config)# route-map GlobalMap permit seq 11
switch(config-route-map-GlobalMap=11)# set weight 100
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100	`config-route-map`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325 8325H 8325P 8360 9300 9300S 10000		

Show commands

show ip aspath-list

```
show ip aspath-list [<NAME>] [vsx-peer]
```

Description

Shows the configuration details of the AS path list.

Parameter	Description
<NAME>	Specifies name of the AS path list.
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Example

Showing the IP AS path list configuration information:

```
switch# show ip aspath-list
ip aspath-list ASLst
    seq 10 permit 22 33
    seq 20 deny 44
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100	Manager (#)	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325 8325H 8325P 8360 9300 9300S 10000		

show ip community-list

```
show ip community-list [<NAME>] [vsx-peer]
```

Description

Shows the configuration details of the community-list.

Parameter	Description
<NAME>	Specifies name of the community-list.
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Example

Showing the community list configuration information:

```
switch# show ip community-list
ip community-list standard CommLst
    seq 10 permit 11:101
    seq 20 deny 12:201
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P	Manager (#)	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8360 9300 9300S 10000		

show ip prefix-list

```
show ip prefix-list [<NAME>] [vsx-peer]
```

Description

Shows the configuration details of the IP prefix lists.

Parameter	Description
<NAME>	Specifies name of the IP prefix lists.
vsx-peer	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Example

Showing the IP prefix list configuration information:

```
switch# show ip prefix-list
ip prefix-list PfxLst: 2 entries
   seq 10 permit 3.0.0.0/8 ge 8 le 8
   seq 20 deny 4.0.0.0/8 ge 8 le 8
```

```
switch# show ipv6 prefix-list
ipv6 prefix-list x: 1 entries
                      seq     10 permit 2011::/64ge 64le 64
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100	Manager (#)	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325 8325H 8325P 8360 9300 9300S 10000		

show route-map

```
show route-map [<NAME>] [vsx-peer]
```

Description

Shows the configuration details of the route map.

Parameter	Description
<NAME>	Specifies name of the route map.
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Example

Showing the route map configuration information:

```
switch# show route-map
Route map: InternetFilter
   Seq 10, permit,
     Match :
       origin : egp
       metric : 123
     Set :
       community : 23:34
       metric : 3
       as_path_exclude : 123
       local_preference : 3456
       origin : igp
       weight : 25
   Seq 20, permit,
     Match :
       origin : egp
       metric : 456
     Set :
       community : 44:44
       metric : 5
       as_path_prepend : 444
       local_preference : 66
       origin : igp
       weight : 250

Route map: LocalFilter
   Seq 10, permit,
```

```
       origin : egp
       metric : 10
                                                    ip next-hop address         : 2.2.2.3
       local-preference            : 20
       route-type                  : external_type1
       source-protocol             : static
       prefix-list                 : PfxLst
       aspath-list                 : ASLst
       community-list              : CommLst
       ip next-hop prefix-list     : PfxLst
       ip route-source prefix-list : PfxLst
     Set :
       community : 22:33
       metric : 25
       as_path prepend : 65535 65534
                                                    ip next-hop address : 2.2.2.4
       local_preference : 30
       origin : igp
       weight : 30
       dampening                   : half-life = 5, reuse = 50, suppress = 125,
max-suppress-time = 15
```

All the match clauses are grouped. All the set clauses are grouped.

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Manager (#)	Administrators or local user group members with execution rights for this command.

Chapter 9

Equal Cost Multipath (ECMP)

Equal Cost Multipath (ECMP) allows a router to load balance traffic destined to some network that is reachable through multiple equal cost route nexthops. ECMP functionality is dependent on L3 routing being enabled on the router.

Overview

The nexthop chosen out of the ECMP group of nexthops for a given packet is typically determined by its header information, usually by hashing source and destination IP addresses as well as potentially source and destination L4 ports. The hashing algorithm used is dependent on the type of switch in use, and user configuration as well as whether the packet is flowing through the slow-path or fast-path.

By default all of the following fields are enabled and utilized for the load balancing decision.

- L3 Destination IP address
- L3 Source IP address
- L4 Destination port
- L4 Source port

ECMP support is NOT dependent on address family. The ECMP feature applies to both IPv4 and IPv6 traffic flowing through the router.

If two different routing protocols are (mis)configured to use the same Administrative Distance, then they have the opportunity to create ECMP groups together.

Dynamic load balancing (DLB) ECMP Group load balancing modes

DLB (dynamic load balancing) ECMP Groups take the amount of bandwidth utilized on each ECMP member into account to assign a packet to an ECMP member.

There are two types of dynamic load balancing modes

- Flow based
- Packet based

Each load balancing mode has different load balancing mechanisms. DLB ECMP Groups only support members that are route-only ports. VLAN interfaces are not supported as members of a DLB Group. ECMP Groups that have any VLAN interface members will be configured to use DLB, and they continue to use static load balancing when either of the DLB mode is enabled.

Flow based dynamic mode

A flow is defined as a stream of packets having the same source IP, destination IP address, source transport port, destination transport port, and same IP protocol. In the flow based dynamic load balance mode, a flow is assigned to an ECMP member which has the highest quality link. Highest quality link is

the link with the least utilization. All ingress packets belonging to a flow remain assigned to the same ECMP member regardless of bandwidth utilization of the other member. Packets of a unique flow are assigned to an ECMP member while they remain active. New flows are assigned to a member based on the current link bandwidth utilization. If a packet of a specific flow is not observed for the duration of the configured inactivity-timeout, then the flow is assigned to a new member the next time it is seen. Flow based dynamic load balance mode avoids reordering of packets at the destination, as long as the configured inactivity-timeout exceeds the maximum path latency within the network. Flow based load balancing is only supported on ports of at least 10 Gbps speed.

Link utilization measurement mechanism

The current link bandwidth utilization calculation is not an instantaneous measure, but instead an exponentially weighted moving average based on a sample rate. If a link has no egress traffic, for example if all traffic is stopped on that link, then the user may expect a new flow to be assigned to this link. However, the instantaneous link utilization may be different than the average link utilization, and so the flow may still be assigned to a different link. If there is more than 1 link of the same quality, random assignment is performed. Overall this provides a more useful metric of link utilization, as the instantaneous view of member imbalance often fluctuates significantly.

Packet based dynamic mode

In the packet based dynamic load balance mode, each ingress packet receives a new ECMP member assignment. ECMP member assignment always uses member with least link utilization for packets belonging to a specific flow is split to different members. This dynamic load balance mode spreads traffic across all ECMP members thus it avoids congestion and provides better network utilization. Packets belonging to a flow can become out of order at the receiving side, and in such case packets need to be reordered.

ECMP commands

ip ecmp dynamic-mode

```
ip ecmp dynamic-mode {(flow-based [inactivity-timeout <256-16384>] | packet-based) outer-header }
```

```
[no] ip ecmp dynamic-mode {(flow-based [inactivity-timeout <256-16384>] | packet-based) outer-header }
```

Description

Specifies the ECMP dynamic mode and its parameters to be used while doing ECMP dynamic load balancing. The no form of this command disables dynamic mode load balancing.

Using the no form of DLB mode that is not currently configured, for example **no ip ecmp dynamic-mode flow-based outer-header** while packet based load balancing is enabled has no effect. Finally, the shortened **no ip ecmp dynamic-mode** command unconfigures _either_ DLB mode, and if no DLB load balancing mode is configured the command has no effect.

Parameter	Description
<inactivity-timeout>	Specifies the inactivity-timeout duration range (Default: 256 microseconds)
<outer-header>	Specifies the Dynamic load-balancing based on outer-header.

Examples

Specifying flow-based dynamic mode with inactivity time of 512 microseconds and dynamic load-balancing based on outer-header:

```
switch(config)# ip ecmp dynamic-mode flow-based inactivity-timeout 512 outer-
header
Modifying ECMP load balancing may cause a brief period of packet loss and
disrupt the network.
Continue (y/n)?
```

Specifying flow-based dynamic mode with the default inactivity-timeout and dynamic load-balancing based on outer-header:

```
switch(config)# ip ecmp dynamic-mode flow-based outer-header
Modifying ECMP load balancing may cause a brief period of packet loss and
disrupt the network.
Continue (y/n)?
```

Specifying packet-based dynamic mode and dynamic load-balancing based on outer-header:

```
switch(config)# ip ecmp dynamic-mode packet-based outer-header
Modifying ECMP load balancing may cause a brief period of packet loss and
disrupt the network.
Continue (y/n)?
```

Removing flow-based dynamic mode:

```
switch(config)# no ip ecmp dynamic-mode flow-based outer-header
Modifying ECMP load balancing may cause a brief period of packet loss and
disrupt the network.
Continue (y/n)?
```

Removing packet-based dynamic mode:

```
switch(config)# no ip ecmp dynamic-mode packet-based outer-header
Modifying ECMP load balancing may cause a brief period of packet loss and
disrupt the network.
Continue (y/n)?
```

Command History

Release	Modification
10.14.1000	Command introduced

Command Information

Platforms	Command context	Authority
9300 9300S	`config`	Administrators or local user group members with execution rights for this command.

show ip ecmp

```
show ip ecmp [vsx-peer]
```

Description

Displays the Equal Cost Multipath (ECMP) configuration.

Parameter	Description
vsx-peer	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Examples

Showing output for HPE Aruba Networking 5420, 6200, 6300,6400, 8320,8325 and 8360 Switch series:

```
switch# show ip ecmp

ECMP Configuration
---------------------

 ECMP Status        : Enabled

ECMP Load Balancing by
------------------------
 Source IP          : Enabled
 Destination IP     : Enabled
 Source Port        : Enabled
 Destination Port   : Enabled
```

Showing output forHPE Aruba Networking 9300, 9300S switch Series:

Flow-based with the default inactivity-timeout:

```
switch# show ip ecmp

ECMP Configuration
---------------------

ECMP Status        : Enabled

Dynamic Load Balancing Mode
----------------------
 Mode               : Flow-based
 Inactivity-timeout : 256 us

ECMP Load Balancing by
------------------------
 Source IP          : Enabled
 Destination IP     : Enabled
 Source Port        : Enabled
 Destination Port   : Enabled
```

Packet-based with the default inactivity-timeout:

```
switch# show ip ecmp

ECMP Configuration
---------------------

ECMP Status        : Enabled

Dynamic Load Balancing Mode
----------------------
Mode               : Packet-based

ECMP Load Balancing by
------------------------
 Source IP          : Enabled
 Destination IP     : Enabled
 Source Port        : Enabled
 Destination Port   : Enabled
```

Dynamic load-balancing has been removed:

```
switch# show ip ecmp

ECMP Configuration
---------------------

ECMP Status        : Enabled

Dynamic Load Balancing Mode
----------------------
Mode               : None

ECMP Load Balancing by
------------------------
Source IP          : Enabled
Destination IP     : Enabled
Source Port        : Enabled
Destination Port   : Enabled
```

Command History

Release	Modification
10.14.1000	Command supported on 9300 and 9300S Switch series.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

Chapter 10

Virtual Router Redundancy Protocol

This section describes how to configure and verify VRRP configuration and operational states.

- All VRRP configurations work in VRRP context.
- VRRP can be configured on physical ports, VLAN interfaces, and LAG interfaces.
- VRRP mandates the associated interface to be routing interface.

Overview

In many networks, edge devices are configured to send packets to a statically configured default router. If this router becomes unavailable, devices that use it as a first-hop router become isolated from the network.

VRRP uses dynamic failover to ensure the availability of an end node default router. The IP address used as the default route is assigned to a virtual router (VR). The VR includes:

- An Active router assigned to forward traffic designated for the virtual route
- One or more prioritized Standby routers (If a Standby is forwarding traffic for the VR, it has replaced the Owner as the Active router.)

This redundancy provides a backup for gateway IP addresses (first-hop routers). If a VR Active router becomes unavailable, the traffic it supports will be transferred to a Standby router without major delays or operator intervention. This operation can eliminate single-point-of-failure problems and provide dynamic failover (and failback) support. As long as one physical router in a VR configuration is available, IP addresses assigned to the VR are always available. Edge devices can send packets to these IP addresses without interruption.

Advantages to using VRRP include:

- Minimizing failover time and bandwidth overhead if a primary router becomes unavailable
- Minimizing service disruptions during a failover
- Providing backup for a load-balanced routing solution
- Avoiding the need to make configuration changes in the end nodes if a gateway router fails
- Eliminating the need for router discovery protocols to support failover operation

Terminology

Owner router: The Owner router is automatically configured with the highest VRRP router priority in the VR (255). It operates as the Active router for the VR unless it becomes unavailable to the network. All Virtual Router members can be configured so virtual IP is not the same as physical (or real) IP. Such virtual address can be called pure virtual IP address.

Active router: The Owner or Standby router that is the physical forwarding agent for routed traffic using the VR as a gateway. There can be only one router operating as the Active for a network. If the router configured as the Owner for a VR is available to the network, it will also be the Active. If the Owner fails or loses availability to the network, the highest-priority Standby becomes the Active.

Standby: A router configured in a VR as a Standby to the Owner/Active for the same VR. There must be a minimum of one Standby in a VR to support VRRP operation if the Owner/Active fails. Every Standby is created with a configurable priority (default: 100). This priority determines the precedence for becoming the Active of the VR if the Owner or another Standby operating as Active becomes unavailable.

When the current Active router is no longer available, the Standby router with highest priority will become the current Active. When a router with higher priority becomes available, it is switched to Active. The user can override this behavior by disabling preemption mode.

Virtual Router (VR): Consists of one Owner/Active router and one or more Standby routers that belong to the same network. The Owner is the router that owns any IP addresses associated with the VR. The VR has one virtual IP address (multiple virtual IP addresses for multinetted interfaces) that correspond to a real IP address on the Owner. A VR includes the following:

- A VR identification (VRID) configured on all VRRP routers in the same network. For a multinetted interface, the VRID is configured on all routers in the same subnet.
- The same VIP configured on each instance of the same VR.
- The status of either Owner or Standby configured on each instance of the same VR. There can be one Owner and one or more Standbys configured on each VR.
- The priority level configured on each instance of the VR. On the Owner router, the highest priority setting of 255 is automatically fixed. On Standbys, the default priority setting is 100 and configurable.
- A VR MAC address that is not configurable.

Virtual IP address (VIP): The VIP associated with a VR must be a real IP address already configured in the associated interface on the Owner router.

- If the VIP is an IPv6 address, a link-local address must be configured before adding a global IPv6 address.
- The Owner and all Standby routers belonging to the VR have this IP address configured in their VRID contexts as the VIP.
- A subnetted interface allows multiple VIPs. If there are fewer IP addresses in an interface and a user wants VRRP support on multiple subnets, configure a separate VR instance for each IP address in the interface.

VRID: The identifier for a VR configured on an interface. A VRID can be used for only one VR in any interface on a router. It can be used again for a different VR in a different interface.

VRRP operation

VRRP supports router redundancy through a prioritized election process among routers configured as members of the same virtual router (VR).

A VR includes two or more member routers configured with:

- A virtual IP address that is also configured as a real IP address on one of the routers
- A virtual router MAC address

The router that owns the IP address is configured to operate as the Owner of the VR for traffic-forwarding purposes. By default, this router has the highest VRRP priority in the VR. Any other routers in the VR have a lower priority and are configured to operate as Standbys if the Owner router becomes unavailable.

The Owner normally operates as the Active for a VR. If the Owner becomes unavailable, a failover to a Standby router belonging to the same VR occurs. This Standby becomes the Active. If the Owner recovers, a failback occurs, and Active status reverts to the Owner.

- Using more than one Standby provides additional redundancy. If the Owner and the highest-priority Standby fail, another lower-priority Standby can take over as Active.
- In a VRRP owner scenario, duplicate address detection (DAD) should be disabled for VRRP to work.

The image shown here depicts a basic VLAN topology. In this example, Routers 1, 2 and 3 form a VRRP group. The IP address of the group is the same address configured for the SVI interface of R1 (10.10.100.1).

Figure 1 *Basic VLAN Topology*

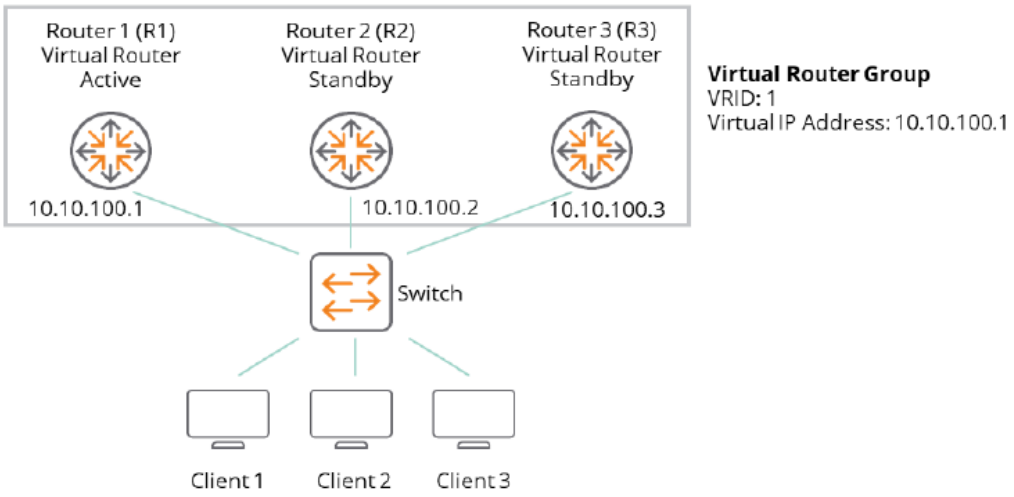

Because the virtual IP address uses the IP address of the SVI interface of Router R1, Router R1 is the Active (also known as IP address Owner).

- As the Active, R1 owns the virtual IP address of the VRRP group and forwards packets sent to this IP address.
- Clients 1 through 3 are configured with the default gateway IP address of 10.10.100.1.
- Routers R2 and R3 function as Standbys. If the Active fails, the Standby router with highest priority becomes the Active and takes over the virtual IP address to provide uninterrupted service for the clients connected to the L2 Switch. When Router R1 recovers, it becomes the Active router again.

Multiple VRRP groups

A user can configure multiple VRRP groups on a router (or L3) interface (Physical, LAG, or SVI interfaces). In a topology in which multiple VRRP groups are configured on a router interface, the interface can act as an Active for one VRRP group and a Standby for one or more other VRRP groups.

This image displays a topology in which VRRP is configured so that Routers R1 and R2 share the traffic to and from clients 1 through 4. Routers R1 and R2 act as Standbys to each other if either router fails.

Figure 1 *Routers share the traffic to and from clients 1 through 4.*

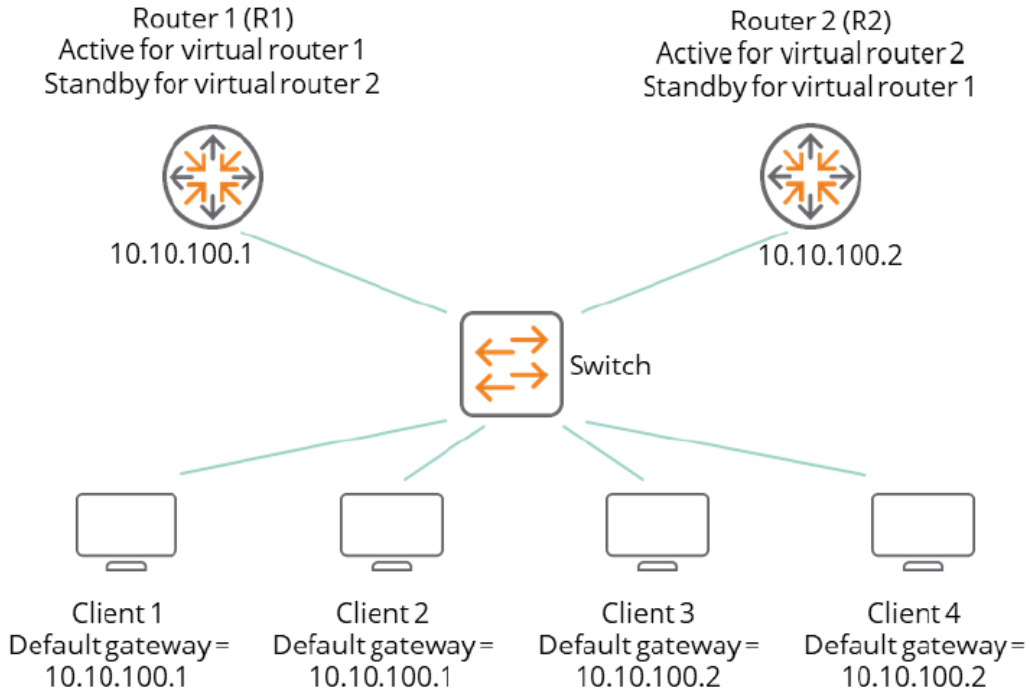

This topology contains two virtual IP addresses for two VRRP groups that overlap. For VRRP group 1, Router R1 is the owner of IP address 10.10.100.1 and is the Active. Router R2 is the Standby to Router R1. Clients 1 and 2 are configured with the default gateway IP address of 10.10.100.1.

For VRRP group 2, Router R2 is the owner of IP address 10.10.100.2 and is the Active. Router R1 is the Standby to Router R2. Clients 3 and 4 are configured with the default gateway IP address of 10.0.0.2.

VRRP priority

In a Standby router VR configuration, the virtual router priority defaults to 100. The priority for the configured Owner is automatically set to the highest value of 255.

In a VR in which there are two or more Standby routers, priority settings can be reconfigured to define the order in which Standbys will be reassigned as Active if a failover occurs.

VRRP preemption

In a Standby router VR configuration, the virtual router priority defaults to 100. The priority for the configured Owner is automatically set to the highest value of 255.

When multiple Standby routers exist in a VR:

- If the current Active fails and the highest-priority Standby is not available, VRRP selects the next-highest priority Standby to operate as Active.
- If the highest-priority Standby later becomes available, it preempts the lower-priority Standby and takes over the Active function.

To remove this preemptive ability on a VR, disable it with the `no preempt` command.

Preemption applies only to VRRP routers configured as Standbys.

Virtual Router MAC address

When a Virtual Router (VR) instance is configured, the protocol automatically assigns a MAC address based on the standard MAC prefix for VRRP packets plus the VRID number (as described in RFC 3768). The first five octets form the standard MAC prefix for VRRP and the last octet is the configured VRID. For example: `00-00-5E-00-01-<VRid>`

VRRP and ARP

The current Active for a VR responds to ARP requests for the virtual IP addresses with the VR-assigned MAC address. The virtual MAC address is also used as source MAC address for periodic advertisements sent by the current Active.

The VRRP router responds to ARP requests for non-virtual IP addresses with the system MAC address. Non-virtual IP addresses are not configured as virtual IP addresses for any VR on the interface.

VRRP and MCLAG

Users can enable VRRP on SVI interfaces (or Layer 3 VLANs) that have multi-chassis lag (MCLAG) ports as member ports. MCLAG allows links that are physically connected to two different AOS-CX devices to appear as a single lag port.

A router performs a Layer 3 route table lookup and Layer 3 forwarding when the destination MAC in the Ethernet frame matches its own MAC address (or virtual MAC address when the router is acting as a VRRP Active). Otherwise the packets are switched.

In a topology with MCLAG and VRRP enabled, an MCLAG switch acting as a VRRP Standby could receive IP packets with the virtual MAC address as the destination MAC. In this scenario, the MCLAG switch acting as a VRRP Standby switches the traffic to the peer (or VRRP Active) using inter-switch-link and the peer (VRRP Active) performs the Layer 3 forwarding.

VRRP tracking

VRRP supports object tracking. This functionality allows monitoring of the state of a configured object. The state determines the priority of the VRRP router in a VRRP group. A track object can be created using the `track <OBJECT-ID>` command and the object can be associated with an interface. A change in interface state will then affect the priority of a VRRP group.

High availability

VRRP supports high availability through stateful restarts and stateful switchovers.

- A stateful restart occurs when the VRRP process (or daemon) fails and is restarted.
- A stateful switchover occurs when the active management module (AMM) switches to the standby management module (SMM).

VRRP and Neighbor Discovery for IPv6

Neighbor Discovery (ND) is the IPv6 equivalent of the IPv4 ARP for layer 2 address resolution. It uses IPv6 ICMP messages to do the following:

- Determine the link-layer address of neighbors
- Verify that a neighbor is reachable
- Track neighbor (local) routers

Neighbor Discovery enables functions such as:

- Router and neighbor solicitations and discovery
- Detecting address changes for devices
- Identifying a replacement for a router or router path that has become unavailable
- Duplicate address detection (DAD)
- Router Advertisement processing
- Neighbor reachability
- Resolution of destination addresses
- Changes to link-layer addresses

An instance of Neighbor Discovery is triggered on a device when a new or changed IPv6 address is detected. VRRPv3 provides a faster failover to a Standby router by not using standard ND procedures. A failover to a Standby router can occur in approximately three seconds without any interaction with hosts and a minimum of VRRPv3 traffic.

Duplicate address detection (DAD)

Duplicate address detection verifies that a configured unicast IPv6 address is unique before it is assigned to an interface. When the Owner router fails, the Standby VRRP router assumes the Active role. When the Owner router becomes operational, DAD will fail because a Standby VRRP router is in the Active role that responds to the DAD request. To avoid this, on virtual routers that are in Owner mode (priority = 255) DAD needs to be disabled for the interface on which the Owner VR is configured.

Guidelines and limitations

- VRRP must be enabled at the global level using the `router vrrp enable` command. It must also be enabled at the interface level at which you configure VRRP.
- A maximum of 8 IPv4 and 8 IPv6 VRRP groups are supported on an interface.
- A maximum of 256 VRRP groups is supported on a router. The groups can be IPv4 or IPv6.
- Supported virtual router identification (VRID) range is 1-255.
- Proxy-ARP and Active-Gateway are not supported in conjunction with VRRP.

VRRP commands

address

```
address <IP-ADDR> [ primary | secondary ]
no address <IP-ADDR> [ primary | secondary ]
```

Description

Configures a primary or secondary IPv4 or IPv6 address for the VRRP group. To use secondary IP addresses in a VRRP group, you must first configure a primary IP address on the same group. A maximum of 16 IP addresses per IPv4 VRRP group and 8 IPv6 addresses per IPv6 VRRP group are

supported.

Do not configure an IPv4 VRRP group using addresses from the /30, /31, and /32 subnets of the interface IP address.

16 Virtual IP addresses per IPv4 VR and 8 Virtual IP addresses per IPv6 VR are supported.

The total number of VIPs supported by a switch is:

- 1024 VIPs for IPv4 VRs
- 512 VIPs for IPv6 VRs

The **no** form of this command deletes a primary or secondary IPv4 or IPv6 address from the VRRP group.

Parameter	Description
<IP-ADDR>	Configures the IPv4 or IPv6 address.
`primary`	Configures a primary address.
`secondary`	Configures a secondary address.

Examples

```
switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# address 10.0.0.1 primary


switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ipv6
switch(config-if-vrrp)# address fe80::1 primary


switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# no address 10.0.0.1 primary


switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ipv6
switch(config-if-vrrp)# no address fe80::1 primary
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

authentication

```
authentication {text | md5} [{plaintext | ciphertext} <KEY>]
no authentication
```

Description

This command enables authentication mode and the authentication key for VRRP groups on an interface. VRRP members or routers of the same VRRP group must use the same authentication mode and authentication key.

The no form of this command disables authentication mode and the authentication key for VRRP groups on an interface.

IPv4 VRRPv3 and IPv6 VRRPv3 do not support VRRP packet authentication. Authentication mode and key configuration take effect only in VRRPv2 (IPv4 only - RFC2338).

In VRRPv3, authentication mode and authentication key settings do not take effect because VRRP Authentication was removed from RFC5798.

VRRP provides the following authentication modes as described in RFC2338:

Simple authentication

The sender fills an authentication key into the VRRP packet and the receiver compares the received authentication key with its local authentication key.

If the two authentication keys match, the received VRRP packet is legitimate. Otherwise, the received packet is illegitimate and is discarded.

Authentication key text is sent in the clear and can be seen in a packet trace. This makes MD5 authentication more secure than text.

MD5 authentication

The sender computes a digest for the packet that will be sent using the authentication key and MD5 algorithm, and saves the result in the VRRP packet.

The receiver performs the same operation with the authentication key and MD5 algorithm, and compares the result with the content in the authentication header. If the results match, the received VRRP packet is legitimate. Otherwise, the received packet is illegitimate and is discarded.

Parameter	Description
`text`	Configures the simple authentication type.

Parameter	Description
`md5`	Configures the MD5 (message-digest) authentication type.
`plaintext`	Specifies that the key is provided as plaintext.
`ciphertext`	Specifies that the key is provided as ciphertext.
<KEY>	Specifies the key in the chosen format.

When the key is not provided on the command line, plaintext key prompting occurs upon pressing Enter. The entered key characters are masked with asterisks.

Examples

Enabling VRRP authentication using MD5 with a provided plaintext key:

```
switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# version 2
switch(config-if-vrrp)# authentication md5 plaintext testvrrpkey
```

Enabling VRRP authentication using MD5 with a prompted plaintext key:

```
switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# version 2
switch(config-if-vrrp)# authentication md5
Enter the authentication key: *************
Re-Enter the authentication key: *************
```

Enabling VRRP authentication using MD5 with a provided ciphertext key:

```
switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# version 2
switch(config-if-vrrp)# authentication md5 ciphertext
AQBapfciFZ/P...biBAAAAOjc0a8=
```

Disabling VRRP authentication:

```
switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# version 2
switch(config-if-vrrp)# no authentication
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if-vrrp`	Administrators or local user group members with execution rights for this command.

bfd *<IPV4-ADDR>*

```
bfd <IPV4-ADDR>
no bfd <IPV4-ADDR>
```

Description

Enables BFD under VRRP for the specified IP address. BFD is asynchronous and echo mode is supported.

The **no** form of this command disables BFD under VRRP for the specified IP address.

Parameter	Description
<IPV4-ADDR>	Specifies the address on which to enable BFD in IPv4 format (**x.x.x.x**), where **x** is a decimal number from 0 to 255.

Examples

On the HPE Aruba Networking 6400 Switch Series, interface identification differs.

Enabling BFD on the address **10.0.0.1** on VRRP **1**:

```
switch(config)# interface 1/1/1
switch(config-if)# routing
switch(config-if)# vrrp 1 address-family ipv4
switch(config-if-vrrp)# bfd 10.0.0.1
```

Disabling BFD on the address **10.0.0.1** on VRRP **1**:

```
switch(config)# interface 1/1/1
switch(config-if)# routing
switch(config-if)# vrrp 1 address-family ipv4
switch(config-if-vrrp)# no bfd 10.0.0.1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if-vrrp`	Administrators or local user group members with execution rights for this command.

preempt

```
preempt
no preempt
```

Description

Enables the preempt option. The default value is enabled. In default mode, a Standby router with a higher priority than another Standby that is operating as Active will take over the Active function.

Applies to VRRP Standby routers only and is used to minimize network disruption caused by unnecessary preemption of the Active operation among Standby routers.

The **no** form of this command disables the preempt option, thus preventing the higher-priority Standby from taking over the Active operation from a lower-priority Standby. This command does not prevent an owner router from resuming the Active function after recovering from being unavailable.

Examples

```
switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# preempt
```

```
switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# no preempt
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config	Administrators or local user group members with execution rights for this command.

preempt delay minimum

```
preempt delay minimum <DELAY-IN-SECONDS>
no preempt delay minimum <DELAY-IN-SECONDS>
```

Description

Sets the time in seconds (1-3600) that the router will wait before taking control of the virtual IP and starting to route packets.

The **no** form of this command sets the preempt delay for the VRRP group to the default preempt delay of 0 seconds.

The VRRP Preempt Delay Timer (PDT) allows admin users to configure a period of time before the VR takes control of the virtual IP address. It does not transition to the Active state until the timer period expires.

The timer value configured should be long enough to allow upper layer protocol to converge. The PDT is applied during initialization and down/up events of the router.

Parameter	Description
<DELAY-IN-SECONDS>	Selects the time in seconds (1-3600). Default is 0 seconds.

Examples

```
switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# preempt delay minimum 30

switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# no preempt delay
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config	Administrators or local user group members with execution rights for this command.

priority

```
priority <1-254>
no priority
```

Description

Sets the priority for the VRRP group.

The **no** form of this command sets the priority for the VRRP group as default priority.

- The default value for non-Owner virtual routers is 100.
- The default value for Owner virtual router is 255, which cannot be changed.

Examples

```
switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# priority 150
```

```
switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# no priority
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360	config	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
9300 9300S 10000		

router vrrp {enable | disable}

```
router vrrp {enable | disable}
no router vrrp {enable | disable}
```

Description

Enables or disables VRRP protocol globally. You must globally enable the VRRP feature for VRRP virtual router.

no router vrrp enable disables VRRP protocol globally but does not remove all VRRP configurations.

no router vrrp disable enables VRRP protocol globally.

Example

Enabling VRRP protocol globally:

```
switch(config)# router vrrp enable
```

Disable VRRP protocol globally:

```
switch(config)# router vrrp disable
```

Disable VRRP protocol globally:

```
switch(config)# no router vrrp enable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config	Administrators or local user group members with execution rights for this command.

no router vrrp

```
no router vrrp
```

Description

Removes VRRP configuration and VRRP global protocol. If **auto-confirm** is enabled or VRRP is not configured on any interface, this command will not ask for user confirmation.

Examples

Removing VRRP configuration:

```
switch(config)# no router vrrp
All VRRP configuration will be deleted.
Do you want to continue (y/n)?
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

show track

```
show track [brief | <OBJECT-ID>]
```

Description

Shows all or specific track object information.

Parameter	Description
`brief`	Displays brief information about all or specific track objects
<OBJECT-ID>	Displays information about a specified track object (1-128)

Examples

```
switch# show track
Track 1
  interface 1/1/1
  Interface is DOWN
```

```
switch# show track brief
Track  Interface       State

1      1/1/1           Down
2      None            Down
3      1/1/2           Up
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Manager (#)	

show track brief

```
show track brief
```

Description

Shows brief information for all track objects.

Examples

Showing brief information for all track objects:

```
switch# show track brief
Track  Interface       State

1      1/1/1           Down
2      None            Down
3      1/1/2           Up
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Manager (#)	

show vrrp

```
show vrrp [brief | detail | interface <INTERFACE-NAME> | interface <LAG-NAME> | interface
<VLAN-NAME> | ip | ipv6 | statistics | statistics interface <INTERFACE-NAME> | statistics
interface <LAG-NAME> | statistics interface <VLAN-NAME>]
[vsx-peer]
```

Description

Shows all VRRP virtual routers information.

Parameter	Description
`brief`	Displays brief output of all VRRP virtual routers Keywords used in displayed information: Grp: VRRP virtual router group ID. A-F: Address Family. Pri: Priority. Time: Uptime of VRRP virtual router since it moved out of INIT state. Pre: Preempt mode (Y is enabled, N if not enabled).
`detail`	Displays detailed output of all VRRP virtual routers
`interface <INTERFACE-NAME>`	Displays VRRP information for a specific interface
`interface <LAG-NAME>`	Displays VRRP information for a specific LAG interface
`interface <VLAN-NAME>`	Displays VRRP information for a specific VLAN interface
`ip`	Displays the IP address family
`ipv6`	Displays IPv6 address family
`statistics`	Displays VRRP statistics information for all interfaces
`statistics interface <INTERFACE-NAME>`	Displays VRRP statistics information for a specific interface

Parameter	Description
`statistics interface <LAG-NAME>`	Displays VRRP statistics information for a specific LAG interface
`statistics interface <VLAN-NAME>`	Displays VRRP statistics information for a specific VLAN interface
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Examples

```
switch# show vrrp
VRRP is enabled

Interface 1/1/1 - Group 1 - Address-Family IPv

  State is ACTIVE
  State duration 56 mins 57.826 secs
  Virtual IP address is 10.0.0.1
  Virtual MAC address is 00:00:5e:00:01:01
  Advertisement interval is 1000 msec
  Preemption enabled
  Priority is 100
  Active Router is 10.0.0.2 (local), priority is 100
  Active Advertisement interval is 1000 msec
  Active Down interval is unknown
  Tracked object ID is 1, and state Down

Interface 1/1/2 - Group 1 - Address-Family IPv4
  State is INIT (Interface Down)
  State duration 45 mins 28.313 secs
  Virtual IP address is no address
  Virtual MAC address is 00:00:5e:00:01:01
  Advertisement interval is 1000 msec
  Preemption enabled
  Priority is 100
  Active Router is unknown, priority is unknown
  Active Advertisement interval is unknown
  Active Down interval is unknown

Interface 1/1/2 - Group 1 - Address-Family IPv6
  State is INIT (Group Disabled)
  State duration 20 mins 19.794 secs
  Virtual IP address is no address
  Virtual secondary IP addresses:
    2201:13::110:4
  Virtual MAC address is 00:00:5e:00:02:01
  Advertisement interval is 1000 msec
  Preemption enabled
  Priority is 100
  Active Router is unknown, priority is unknown
  Active Advertisement interval is unknown
  Active Down interval is unknown
```

```
switch# show vrrp brief

VRRP is enabled

 Interface   Grp  A-F   Pri   Time Owner Pre State   Active addr/Group addr
  1/1/1       1   IPv4  100    0    N    Y   ACTIVE  10.0.0.2(local) 10.0.0.1
  1/1/2       1   IPv4  100    0    N    Y   INIT    AF-UNDEFINED no address
  1/1/2       1   IPv6  100    0    N    Y   INIT    AF-UNDEFINED no address
```

```
switch# show vrrp detail
VRRP is enabled

Interface 1/1/1 - VRRPv2 Statistics
   Invalid group ID packets received : 0
   Invalid version packets received : 0
   Invalid checksum packets received : 0

Interface 1/1/1 - VRRPv3 Statistics
   Invalid group ID packets received : 0
   Invalid version packets received : 0
   Invalid checksum packets received : 0

Interface 1/1/1 - Group 1 - Address-Family IPv4
  State is ACTIVE
  State duration 1 mins 35.486 secs
  Virtual IP address is 10.0.0.1
  Virtual MAC address is 00:00:5e:00:01:01
  Advertisement interval is 1000 msec
  Version 3
  Preemption enabled
  Priority is 100
  Active Router is 10.0.0.2 (local), priority is 100
  Active Advertisement interval is 1000 msec
  Active Down interval is unknown
  Tracked object ID is 1, and state Down
  VRRPv3 Advertisements: sent 3931 (errors 0) - rcvd 0
  VRRPv2 Advertisements: sent 0 (errors 0) - rcvd 0
  Group Discarded Packets: 3537
    IP address Owner conflicts: 0
    IP address configuration mismatch : 3537
    Advert Interval errors : 0
    Adverts received in Init state: 0
    Invalid group other reason: 0
  Group State transition:
    Init to active: 0
    Init to standby: 2 (Last change Mon Jun 16 11:19:36.316 UTC)
    Standby to active: 2 (Last change Mon Jun 16 11:19:39.926 UTC)
    Active to standby: 0
    Active to init: 1 (Last change Mon Jun 16 11:17:49.978 UTC)
    Standby to init: 0

Interface 1/1/2 - VRRPv2 Statistics
   Invalid group ID packets received : 0
   Invalid version packets received : 0
   Invalid checksum packets received : 0

Interface 1/1/2 - VRRPv3 Statistics
   Invalid group ID packets received : 0
   Invalid version packets received : 0
   Invalid checksum packets received : 0
```

```
Interface 1/1/2 - Group 1 - Address-Family IPv4
  State is INIT (Interface Down)
  State duration 49 mins 23.507 secs
  Virtual IP address is no address
  Virtual MAC address is 00:00:5e:00:01:01
  Advertisement interval is 1000 msec
  Version 3
  Preemption enabled
  Priority is 100
  Active Router is unknown, priority is unknown
  Active Advertisement interval is unknown
  Active Down interval is unknown
  VRRPv3 Advertisements: sent 0 (errors 0) - rcvd 0
  VRRPv2 Advertisements: sent 0 (errors 0) - rcvd 0
  Group Discarded Packets: 0
    IP address Owner conflicts: 0
    IP address configuration mismatch : 0
    Advert Interval errors: 0
    Adverts received in Init state: 0
    Invalid group other reason: 0
  Group State transition:
    Init to active: 0
    Init to standby: 0
    Standby to active: 0
    Active to standby: 0
    Active to init: 0
    Standby to init: 0

Interface 1/1/2 - Group 1 - Address-Family IPv6
  State is INIT (Interface Down)
  State duration 24 mins 14.988 secs
  Virtual IP address is no address
  Virtual secondary IP addresses:
    2201:13::110:4
  Virtual MAC address is 00:00:5e:00:02:01
  Advertisement interval is 1000 msec
  Preemption enabled
  Priority is 100
  Active Router is unknown, priority is unknown
  Active Advertisement interval is unknown
  Active Down interval is unknown
  VRRPv3 Advertisements: sent 0 (errors 0) - rcvd 0
  VRRPv2 Advertisements: sent 0 (errors 0) - rcvd 0
  Group Discarded Packets: 0
    VRRPv2 incompatibility: 0
    IP address Owner conflicts: 0
    IP address configuration mismatch : 0
    Advert Interval errors : 0
    Adverts received in Init state: 0
    Invalid group other reason: 0
  Group State transition:
    Init to active: 0
    Init to standby: 0
    Standby to active: 0
    Active to standby: 0
    Active to init: 0
    Standby to init: 0
```

```
switch# show vrrp interface 1/1/1
VRRP is enabled
```

```
Interface 1/1/1 - Group 1 - Address-Family IPv4
  State is ACTIVE
  State duration 11 mins 21.617 secs
  Virtual IP address is 10.0.0.1
  Virtual MAC address is 00:00:5e:00:01:01
  Advertisement interval is 1000 msec
  Version 3
  Preemption enabled
  Priority is 100
  Active Router is 10.0.0.2 (local), priority is 100
  Active Advertisement interval is 1000 msec
  Active Down interval is unknown
```

```
switch# show vrrp interface lag10
VRRP is enabled

Interface lag10 - Group 1 - Address-Family IPv4
  State is ACTIVE
  State duration 11 mins 21.617 secs
  Virtual IP address is 10.0.0.1
  Virtual MAC address is 00:00:5e:00:01:01
  Advertisement interval is 1000 msec
  Version 3
  Preemption enabled
  Priority is 100
  Active Router is 10.0.0.2 (local), priority is 100
  Active Advertisement interval is 1000 msec
  Active Down interval is unknown
```

```
switch# show vrrp interface vlan100
VRRP is enabled

Interface vlan100 - Group 1 - Address-Family IPv4
  State is ACTIVE
  State duration 11 mins 21.617 secs
  Virtual IP address is 10.0.0.1
  Virtual MAC address is 00:00:5e:00:01:01
  Advertisement interval is 1000 msec
  Version 3
  Preemption enabled
  Priority is 100
  Active Router is 10.0.0.2 (local), priority is 100
  Active Advertisement interval is 1000 msec
  Active Down interval is unknown
```

```
switch# show vrrp statistics

VRRP is enabled

Interface 1/1/1 - VRRPv2 Statistics
   Invalid group ID packets received : 0
   Invalid version packets received : 0
   Invalid checksum packets received : 0

Interface 1/1/1 - VRRPv3 Statistics
   Invalid group ID packets received : 0
   Invalid version packets received : 0
   Invalid checksum packets received : 0
```

```
VRRP Statistics for interface 1/1/1 - Group 1 - Address-Family IPv4
  State is ACTIVE
  State duration 6 mins 55.006 secs
  VRRPv3 Advertisements: sent 4288 (errors 0) - rcvd 0
  VRRPv2 Advertisements: sent 0 (errors 0) - rcvd 0
  Group Discarded Packets: 3856
    IP address Owner conflicts: 0
    IP address configuration mismatch : 0
    Advert Interval errors : 0
    Adverts received in Init state: 0
    Invalid group other reason: 0
  Group State transition:
    Init to active: 0
    Init to standby: 2 (Last change Mon Jun 16 11:19:36.316 UTC)
    Standby to active: 2 (Last change Mon Jun 16 11:19:39.926 UTC)
    Active to standby: 0
    Active to init: 1 (Last change Mon Jun 16 11:17:49.978 UTC)
    Standby to init: 0

Interface 1/1/2 - VRRPv2 Statistics
   Invalid group ID packets received : 0
   Invalid version packets received : 0
   Invalid checksum packets received : 0

Interface 1/1/2 - VRRPv3 Statistics
   Invalid group ID packets received : 0
   Invalid version packets received : 0
   Invalid checksum packets received : 0

VRRP Statistics for Interface 1/1/2 - Group 1 - Address-Family IPv4
  State is INIT (No Primary Group Address)
  State duration 54 mins 43.027 secs
  VRRPv3 Advertisements: sent 0 (errors 0) - rcvd 0
  VRRPv2 Advertisements: sent 0 (errors 0) - rcvd 0
  Group Discarded Packets: 0
    IP address Owner conflicts: 0
    Invalid address count: 0
    IP address configuration mismatch : 0
    Advert Interval errors : 0
    Adverts received in Init state: 0
    Invalid group other reason: 0
  Group State transition:
    Init to active: 0
    Init to standby: 0
    Standby to active: 0
    Active to standby: 0
    Active to init: 0
    Standby to init: 0

VRRP Statistics for Interface 1/1/2 - Group 1 - Address-Family IPv6
  State is INIT (Interface Down)
  State duration 29 mins 34.508 secs
  VRRPv3 Advertisements: sent 0 (errors 0) - rcvd 0
  VRRPv2 Advertisements: sent 0 (errors 0) - rcvd 0
  Group Discarded Packets: 0
    IP address Owner conflicts: 0
    IP address configuration mismatch : 0
    Advert Interval errors: 0
    Adverts received in Init state: 0
    Invalid group other reason: 0
  Group State transition:
```

```
    Init to active: 0
    Init to standby: 0
    Standby to active: 0
    Active to standby: 0
    Active to init: 0
    Standby to init: 0
```

```
switch# show vrrp statistics interface 1/1/1

VRRP is enabled

Interface 1/1/1 - VRRPv2 Statistics
   Invalid group ID packets received : 0
   Invalid version packets received : 0
   Invalid checksum packets received : 0

Interface 1/1/1 - VRRPv3 Statistics
   Invalid group ID packets received : 0
   Invalid version packets received : 0
   Invalid checksum packets received : 0

VRRP Statistics for interface 1/1/1 - Group 1 - Address-Family IPv4
  State is ACTIVE
  State duration 6 mins 55.006 secs
  VRRPv3 Advertisements: sent 4288 (errors 0) - rcvd 0
  VRRPv2 Advertisements: sent 0 (errors 0) - rcvd 0
  Group Discarded Packets: 3856
    IP address Owner conflicts: 0
    IP address configuration mismatch : 0
    Advert Interval errors : 0
    Adverts received in Init state: 0
    Invalid group other reason: 0
  Group State transition:
    Init to active: 0
    Init to standby: 2 (Last change Mon Jun 16 11:19:36.316 UTC)
    Standby to active: 2 (Last change Mon Jun 16 11:19:39.926 UTC)
    Active to standby: 0
    Active to init: 1 (Last change Mon Jun 16 11:17:49.978 UTC)
    Standby to init: 0
```

```
switch# show vrrp statistics interface lag10

VRRP is enabled

Interface lag10 - VRRPv2 Statistics
   Invalid group ID packets received : 0
   Invalid version packets received : 0
   Invalid checksum packets received : 0

Interface lag10 - VRRPv3 Statistics
   Invalid group ID packets received : 0
   Invalid version packets received : 0
   Invalid checksum packets received : 0

VRRP Statistics for interface lag10 - Group 1 - Address-Family IPv4
  State is ACTIVE
  State duration 6 mins 55.006 secs
  VRRPv3 Advertisements: sent 4288 (errors 0) - rcvd 0
  VRRPv2 Advertisements: sent 0 (errors 0) - rcvd 0
  Group Discarded Packets: 3856
```

```
    IP address Owner conflicts: 0
    IP address configuration mismatch : 0
    Advert Interval errors : 0
    Adverts received in Init state: 0
    Invalid group other reason: 0
  Group State transition:
    Init to active: 0
    Init to standby: 2 (Last change Mon Jun 16 11:19:36.316 UTC)
    Standby to active: 2 (Last change Mon Jun 16 11:19:39.926 UTC)
    Active to standby: 0
    Active to init: 1 (Last change Mon Jun 16 11:17:49.978 UTC)
    Standby to init: 0
```

```
switch# show vrrp statistics interface vlan100

VRRP is enabled

Interface vlan100 - VRRPv2 Statistics
   Invalid group ID packets received : 0
   Invalid version packets received : 0
   Invalid checksum packets received : 0

Interface vlan100 - VRRPv3 Statistics
   Invalid group ID packets received : 0
   Invalid version packets received : 0
   Invalid checksum packets received : 0

VRRP Statistics for interface vlan100 - Group 1 - Address-Family IPv4
  State is ACTIVE
  State duration 6 mins 55.006 secs
  VRRPv3 Advertisements: sent 4288 (errors 0) - rcvd 0
  VRRPv2 Advertisements: sent 0 (errors 0) - rcvd 0
  Group Discarded Packets: 3856
    IP address Owner conflicts: 0
    IP address configuration mismatch : 0
    Advert Interval errors : 0
    Adverts received in Init state: 0
    Invalid group other reason: 0
  Group State transition:
    Init to active: 0
    Init to standby: 2 (Last change Mon Jun 16 11:19:36.316 UTC)
    Standby to active: 2 (Last change Mon Jun 16 11:19:39.926 UTC)
    Active to standby: 0
    Active to init: 1 (Last change Mon Jun 16 11:17:49.978 UTC)
    Standby to init: 0
```

Command History

Release	Modification
10.14	Replaced the **ipv4** parameter with the **ip** parameter. The **ipv4** parameter is deprecated.
10.08	Updated command output for inclusive language
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Manager (#)	

shutdown

```
shutdown
no shutdown
```

Description

Enables standby VRRP groups on the interface to route the traffic sent to virtual router's MAC address.

It is recommended to enable this mode on VLANs where mc-lags are configured.

Disabled by default.

The **no** form of this command disables standby VRRP groups on the interface to route the traffic sent to virtual router's MAC address.

Only supported on SVI interfaces.

Examples

Enabling standby VRRP groups on interface VLAN 10 to route traffic sent to the router's MAC address:

```
switch(config)# interface vlan 10
switch(config-if-vlan)# vrrp dual-active-forwarding
```

Disabling the ability of standby VRRP groups on interface VLAN 10 to route traffic sent to the router's MAC address:

```
switch(config)# interface vlan 10
switch(config-if-vlan)# no vrrp dual-active-forwarding
```

Command History

Release	Modification
10.12.1000	Command introduced.

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8360 9300 10000 9300S 10000	config	Administrators or local user group members with execution rights for this command.

timers advertise

```
timers advertise <ADVERTISE-IN-MILLISECONDS>
no timers advertise
```

Description

Sets the advertisement interval in ms (100-40950). The default value is 1000. Advertisement interval can be configured in multiples of 1,000 ms.

The **no** form of this command sets the advertisement interval in ms to the default value of 1000.

This release does not support sub-second timer for VRRPv3.

Examples

Setting the advertisement interval in ms to 2000:

```
switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# timers advertise 2000
```

Setting the advertisement interval in ms to the default value of 1000:

```
switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# no timers advertise
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100	config	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325 8325H 8325P 8360 9300 9300S 10000		

track (VRRP group)

```
track <OBJECT-ID>
no track <OBJECT-ID>
```

Description

Sets the track object ID (1-128) for the group. The track object is first configured globally for the interface and then attached to the VRRP virtual router.

The track object must not track the same interface for which a VRRP group is configured.

The **no** form of this command removes the track object ID from the group.

Examples

Setting the track object ID for the group:

```
switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# track 1
```

Removing the track object ID from the group:

```
switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# no track 1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P	config	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8360 9300 9300S 10000		

track (VRRP virtual router)

```
track <OBJECT-ID>
no track <OBJECT-ID>
```

Description

Configures a track object that can be associated with an interface. A change in interface state will then affect the priority of a VRRP group. By default, no interface is associated to a track object, so state is down.

The **no** form of this command deletes a tracked object for an interface. If it is not associated with a VRRP virtual router, a track object cannot be deleted.

Track cannot be configured by using port with no routing.

When all tracked interfaces go down on a virtual router, priority is automatically set to zero instead of its configured value. Owner virtual routers always use a default priority of 255.

Parameter	Description
<OBJECT-ID>	Specify the track object ID value. Range: 1 to 128.

Examples

Configuring a tracked object:

```
switch(config)# track 1
```

Deleting a tracked object:

```
switch(config)# no track 1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400	config	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300 9300S 10000		

track by

```
track by <OBJECT-ID>
no track by <OBJECT-ID>
```

Description

Specifies an interface to be tracked when changes in the state of the interface affect the priority of a VRRP group. Once track is associated with an interface, the track state reflects the interface forwarding state.

The **no** form of this command removes an interface from tracking, affecting VRRP states of any interfaces associated with VRRP groups.

The VLAN interface 1 is always tracked.

Parameter	Description
<OBJECT-ID>	Specifies the track object ID value. Range: 1 to 128.

Example

Specifying an interface to be tracked:

```
switch (config)# interface 1/1/1
switch (config-if)# track by 1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H	config	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325P 8360 9300 9300S 10000		

version

`version <VERSION-NUMBER>`

Description

Sets the protocol version for the VRRP group. Version change is allowed only for the IPv4 address-family. The default value is 2, which supports IPv4 with minimum 1 second advertisement interval. Value 3 supports IPv4 and IPv6 with minimum 1 second advertisement interval.

Parameter	Description
<VERSION-NUMBER>	Specifies the VRRP protocol version. Possible values: 2 or 3. The default value is 2, which supports IPv4 with a minimum 1 second advertisement interval.

Example

Setting the protocol version for the VRRP group to 3:

```
switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)# version 3
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

vrrp

```
vrrp <VRID> address-family {ip | ipv6}
no vrrp <VRID> address-family {ip | ipv6}
```

Description

Creates a VRRP group and establishes VRRP group configuration context.

- A maximum of 16 VRRP groups, including both IPv4 and IPv6, are supported on an interface.
- A maximum of 256 VRRP groups is supported on a router. The groups can be IPv4 or IPv6 on a first come first serve basis.

The **no** form of this command deletes a VRRP group.

Parameter	Description
<VRID>	Selects the VRRP router ID value. Range: 1 to 255.
`address-family [IP| IPv6]`	Specifies which address family to use, IP or IPv6.

Examples

```
switch(config)# interface 1/1/1
switch(config-if)# vrrp 1 address-family ip
switch(config-if-vrrp)#

switch(config-if-vrrp)# no vrrp 1 address-family ip
```

Command History

Release	Modification
10.14	Replaced the **ipv4** parameter with the **ip** parameter. The **ipv4** parameter is deprecated.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

vrrp dual-active-forwarding

```
vrrp dual-active-forwarding
no vrrp dual-active-forwarding
```

Description

Allows standby VRRP groups on an interface to route traffic sent to a virtual router's MAC address.

It is recommended to enable VRRP dual active forwarding on VLANs configured on MCLAG interfaces.

Disabled by default.

The **no** form of this command prevents standby VRRP groups on the interface to route traffic sent to the virtual router's MAC address.

VRRP dual active forwarding is only supported on SVI interfaces.

VRRP dual active forwarding and VRRP owner mode are mutually exclusive.

To verify the configuration of **vrrp dual-active-forwarding** use the `show vrrp` command.

Examples

Enabling VRRP dual active forwarding on interface VLAN 10:

```
switch(config)# interface vlan 10
switch(config-if-vlan)# vrrp dual-active-forwarding
```

Disabling VRRP dual active forwarding on interface VLAN 10:

```
switch(config)# interface vlan 10
switch(config-if-vlan)# no vrrp dual-active-forwarding
```

Command History

Release	Modification
10.12.1000	Command introduced.

Command Information

Platforms	Command context	Authority
6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if-vlan`	Administrators or local user group members with execution rights for this command.

Chapter 11
Inter-Virtual Router Forwarding (IVRF)

Virtual Routing and Forwarding (VRF) allows a physical router to be partitioned into multiple Virtual Router Forwarding instances. The control and data plane are isolated in each VRF. This isolation keeps the traffic from crossing VRFs, and therefore multiple routing tables can coexist within the same physical L3 switch/router.

Route leaking allows for routes to leak (be distributed), across multiple VRFs in a controlled manner. Two methods to leak routes from one VRF to another are:

- Using static routes.
- Using Multiprotocol BGP (MP-BGP) (dynamic route leaking using route-targets).

Troubleshooting IVRL

If the RADIUS server configuration is not in the same VRF as the actual server, then the client authentication can fail because the server times out.In the following configuration, the RADIUS server configuration is in **VRF ADMIN** and the IP is reachable on **VRF USER**. This configuration can cause port-access client authentication to fail due to a server timeout error.

```
radius-server host 10.1.1.20 key ciphertext
AQBapYuQcBB2A0RVgSPc/7H5lfSBrt8rwbNhTouCxjeE4L8SCQAAAGZ3ykq9PEswPA== tracking
enable vrf ADMIN
radius-server host 10.1.1.21 key ciphertext
AQBapYuQcBB2A0RVgSPc/7H5lfSBrt8rwbNhTouCxjeE4L8SCQAAAGZ3ykq9PEswPA== tracking
enable vrf ADMIN
interface 1/3/47
  no shutdown
  routing
  vrf attach USER
  ip address 10.10.100.1/30
ip route 10.1.1.0/24 10.10.100.2 vrf USER
```

Best practices is to either have the RADIUS server configuration in the VRF to which the server is connected, or to leak the connected route to the VRF with the RADIUS server configuration.

This example resolves the issue in the previous configuration by leaking route to the VRF **VRF ADMIN.**

```
p route 10.10.100.0/30 1/1/1 vrf ADMIN"
```

Static VRF route leaking

Static VRF route leaking between two VRFs is allowed using an IP route configuration command. To accomplish the route leak:

- The routes being leaked to other VRFs must be present in the source VRF.
- When a route leak has been accomplished and is therefore present in the source VRF, that route can then be leaked to other VRFs.

Dynamic VRF route leaking

Dynamic VRF route leaking is allowed between two or more VRFs. Dynamic VRF is configured using export and import route-targets. Multiprotocol BGP (MP-BGP) is used to leak routes between the VRFs. To configure MP-BGP, a separate address-family must be configured within BGP for each VRF, and simply redistribute routes (OSPF, static or connected) within that VRF for the routes to be leaked.

Dynamic VRF route leak restrictions and limitations

- The maximum number of route targets supported in each VRF, including in the default VRF context, is 256.
- The maximum number of supported dynamic leaked routes per system is 16K.
- Multicast route leaking is not supported.
- On the same switch, more than one level of Inter-VRF route leaking, or cascaded IVRL, is not supported. For example, VRF1 to VRF2 with the default being the intermediate VRF.
- The EVPN address family is not supported under the default VRF.
- The mgmt VRF cannot be used for VRF rout leaking.
- Extended Community cannot be sent to the neighbors with the prefix.

Procedure to leak routes between VRFs

1. Enter the VRF context (if the named VRF does not exist, it is created).
2. Configure the route distinguisher (RD) for each VRF.
3. Configure the appropriate address-family for each VRF.
4. Configure the BGP route targets.
5. Associate the VRFs with the interfaces.
6. Configure IP addresses on the interface.
7. Enable the required IGP protocol (for example, OSPF) on the VRF.
8. Configure multiprotocol BGP.
9. Redistribute the routes (OSPF, static or connected routes) to be leaked into BGP. Multiprotocol BGP (MP-BGP) will be used to leak routes between VRFs. To configure MP-BGP, you must configure a separate address-family within BGP for each VRF, and simply redistribute routes (OSPF, static or connected) within that VRF for the routes to be leaked.

Troubleshooting inter-VRF route leaking

If the RADIUS server configuration is not in the same VRF as the actual server, client authentication can fail because the server times out. In the following example configuration, the RADIUS server configuration is in **VRF ADMIN** and the IP is reachable on **VRF USER**. This configuration can cause port-access client authentication to fail due to a server timeout error.

```
radius-server host 10.1.1.20 key ciphertext
AQBapYuQcBB2A0RVgSPc/7H5lfSBrt8rwbNhTouCxjeE4L8SCQAAAGZ3ykq9PEswPA== tracking
enable vrf ADMIN
radius-server host 10.1.1.21 key ciphertext
AQBapYuQcBB2A0RVgSPc/7H5lfSBrt8rwbNhTouCxjeE4L8SCQAAAGZ3ykq9PEswPA== tracking
enable vrf ADMIN
interface 1/3/47
  no shutdown
  routing
  vrf attach USER
  ip address 10.10.100.1/30
ip route 10.1.1.0/24 10.10.100.2 vrf USER
```

Best practices is to either have the RADIUS server configuration in the VRF to which the server is connected, or to leak the connected route to the VRF with the RADIUS server configuration.

This example resolves the issue in the previous configuration by leaking the route to the VRF **VRF ADMIN.**

```
ip route 10.10.100.0/30 1/1/1 vrf ADMIN"
```

IVRF commands

address-family

```
address-family [<AFI> | <SAFI>] [ip | ipv6] [unicast]
[no] address-family <AFI> | <SAFI>
```

Description

Initializes the appropriate address-family and enters address-family configuration mode for IPv4 or IPv6. The unicast option is available to configure the subaddress family identifier.

The **no** form of the command removes the association of the specified address-family. The address-family specific routes that are leaked from this VRF will be withdrawn.

Parameter	Description
AFI	Required: Specifies address family identifier.
SAFI	Required: Specifies subaddress family identifier.
ip	Optional: IPv4 address family
ipv6	Optional: IPv6 address family
unicast	The subaddress family identifier. When the **unicast** option is used, the command context changes to **config-vrf-af-ipv4-uc**.

Examples

Address family command for IPv4 unicast:

```
switch(config-vrf)# address-family ipv4 unicast
switch(config-vrf)#no address-family ipv4 unicast
```

Address family command for IPv6 unicast:

```
switch(config-vrf)# address-family ipv6 unicast
switch(config-vrf)#no address-family ipv6 unicast
```

Command History

Release	Modification
10.14	The **ipv4** keyword is deprecated and replaced with **ip**.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-vrf	Administrators or local user group members with execution rights for this command.

ip|ipv6 vrf

```
[ip|ipv6] route <PREFIX> <SRC-VRF-LOCAL-IFACE><SRC-VRF-NEXTHOP-IP> vrf <DST-VRF-NAME>
[name <WORD>]
no [ip|ipv6] route <PREFIX> <SRC-VRF-LOCAL-IFACE><SRC-VRF-NEXTHOP-IP> vrf <DST-VRF-NAME>
[name <WORD>]
```

Description

The IP/IPv6 route command sets the subnet mask, the reachable network interface, the next-hop IP for the reachable network, and the VRF route leak destination.

Parameter	Description
<PREFIX>	The subnet mask (prefix of the network).
<SRC-VRF-LOCAL-IFACE>	The interface which is reachable by the network.
<SRC-VRF-NEXTHOP-IP>	The next-hop IP for the reachable network.
<DST-VRF-NAME>	The VRF route leak destination.
name *<WORD>*	Assigns a name to the destination route. Range: 1 to 64 characters.

Examples

Using the command, leak the named route Blue VRF, using prefix 100.0.0.0/24 which is reachable by the next-hop IP 20.0.0.1 on the interface 1/1/1 from VRF Red:

```
switch(config)# ip route 100.0.0.0/24 1/1/1 20.0.0.1 vrf blue
```

Assign a name to the destination route:

```
switch(config)# ip route 20.0.4.0/24 10.0.0.2 vrf test name VRF-Management
```

The following example shows the route in the source VRF, the corresponding command executed in the other VRF, and howidcates how the routing tables is displayed after the command is executed.

```
switch(config)# show running-config
Current configuration:
!
vrf blue
vrf green
vrf red
!
vlan 1
interface 1/1/19
   no shutdown
   vrf attach red
   ip address 10.0.0.1/24
interface 1/1/20
   no shutdown
   vrf attach green
   ip address 20.0.0.1/24
interface 1/1/21
   no shutdown
   vrf attach blue
   ip address 40.0.0.1/24
interface 1/1/22
no shutdown

switch(config)# ip route 30.0.0.0/24 20.0.0.2 vrf green
switch(config)# ip route 50.0.0.0/24 20.0.0.2 vrf green
switch(config)# ip route 30.0.0.0/24 1/1/20 20.0.0.2 vrf red
switch(config)# ip route 50.0.0.0/24 1/1/20 20.0.0.2 vrf blue
switch(config)# ip route 40.0.0.0/24 1/1/21 vrf red


switch# show ip route vrf green
Displaying ipv4 routes selected for forwarding
Origin Codes: C - connected, S - static, L - local
R - RIP, B - BGP, O - OSPF, D - DHCP
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
IA - OSPF internal area, E1 - OSPF external type 1
E2 - OSPF external type 2
Flags:        F - FIB-optimized route
VRF: green
Prefix              Nexthop                                  Interface     VRF
(egress)       Origin/   Distance/    Age
Type      Metric
-------------------------------------------------------------------------------
----------------------
20.0.0.0/24         -                                        1/1/20        -
```

```
          C         [0/0]       -
20.0.0.1/32         -                                       1/1/20       -
          L         [0/0]       -
30.0.0.0/24         20.0.0.2                                1/1/20       -
          S         [1/0]       00h:05m:16s
50.0.0.0/24         20.0.0.2                                1/1/20       -
          S         [1/0]       00h:05m:16s
Total Route Count : 4


switch# show ip route vrf blue
Displaying ipv4 routes selected for forwarding
Origin Codes: C - connected, S - static, L - local
R - RIP, B - BGP, O - OSPF, D - DHCP
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
IA - OSPF internal area, E1 - OSPF external type 1
E2 - OSPF external type 2
Flags:        F - FIB-optimized route
VRF: blue
Prefix              Nexthop                                 Interface     VRF
(egress)       Origin/   Distance/    Age
Type      Metric
-------------------------------------------------------------------------------
----------------------
40.0.0.0/24         -                                       1/1/21       -
          C         [0/0]       -
40.0.0.1/32         -                                       1/1/21       -
          L         [0/0]       -
50.0.0.0/24         20.0.0.2                                1/1/20       green
          S         [1/0]       00h:00m:44s
Total Route Count : 3


switch# show ip route vrf red
Displaying ipv4 routes selected for forwarding
Origin Codes: C - connected, S - static, L - local
R - RIP, B - BGP, O - OSPF, D - DHCP
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
IA - OSPF internal area, E1 - OSPF external type 1
E2 - OSPF external type 2
Flags:        F - FIB-optimized route
VRF: red
Prefix              Nexthop                                 Interface     VRF
(egress)       Origin/   Distance/    Age
Type      Metric
-------------------------------------------------------------------------------
----------------------
10.0.0.0/24         -                                       1/1/19       -
          C         [0/0]       -
10.0.0.1/32         -                                       1/1/19       -
          L         [0/0]       -
30.0.0.0/24         20.0.0.2                                1/1/20       green
          S         [1/0]       00h:00m:48s
40.0.0.0/24         -                                       1/1/21       blue
          S         [1/0]       00h:00m:48s
100.100.1.0/24      -                                       vlan10       -
          C         [0/0]       -
100.100.1.7/32      -                                       vlan10       -
          L         [0/0]       -
Total Route Count : 6
```

Command History

Release	Modification
10.17	The **name** parameter was introduced.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

ipv6 route source interface

```
ipv6 route <IPv6 ADDR/MEMBER><IPv6 ADDR> source  <INTERFACE>
 [<DISTANCE> | <VRF INSTANCE-NAME>]
no ipv6 route
```

Description

Creates a route leak between the source VRF and destination VRF. Using the static method, the route must first be added to the destination VRF. The route is added to the local interface of the source VRF with a next-hop interface. The existing IPv6 route command takes the source interface only when next-hop IP is link-local. To support VRF route leaking for global IPv6 unicast addresses, the command takes next-hop interface information along with next-hop IP regardless of next-hop IP is link-local or not.

- Users must provide both the next-hop IP and the interface information to leak the global unicast IPv6 network routes (route that is not directly reachable).
- The next-hop IP information is not required to leak connected global unicast IPv6 routes (route that is directly reachable).

The **no** form of command deletes the static VRF leaked route.

Parameter	Description
`<IPv6 ADDR/MEMBER>`	Required: IPv6 IP-Address route destination.
`<IPv6 ADDR>`	Required: IPv6 route destination
`<INTERFACE>`	Required: The outgoing interface. Use the format `member/slot/port` (for example, `1/3/1`).
`<DISTANCE>`	Optional: administrative distance of static route
`<VRF INSTANCE-NAME>`	Optional: VRF instance

Options

`nullroute`

Discard packets to the destined route silently.

`reject`

Discard packets to the destined route and return ICMP error to the sender.

Examples

Configures a route leak between the source VRF and destination VRF:

```
switch(config)# show runn
Current configuration:
!
vrf blue
      vrf green
vrf red
      !

vlan 1
               interface 1/1/1
               no shutdown
                      vrf attach red
        ip address 2000::1/64
                      interface 1/1/2
                      no shutdown
                      vrf attach green
               ip address 3000::1/64
               interface 1/1/3
                      no shutdown
               vrf attach blue
               ip address 4000::1/64


      switch(config)# ipv6 route 5000::0/64 3000::2 source 1/1/2 vrf red
switch(config)# ipv6 route 6000::0/64 3000::3 source 1/1/2 vrf blue


      switch(config)# show runn
Current configuration:
!
vrf blue
      vrf green
vrf red
      !

vlan 1
               interface 1/1/1
               no shutdown
                      vrf attach red
               ip address 2000::1/64
                      interface 1/1/2
                      no shutdown
                      vrf attach green
               ip address 3000::1/64
               interface 1/1/3
                      no shutdown
               vrf attach blue
        ip address 4000::1/64

        ipv6 route 5000::0/64 3000::2 source 1/1/2 vrf red
```

```
  ipv6 route 6000::0/64 3000::3 source 1/1/2 vrf blue

              switch(config)# no ipv6 route 5000::0/64 3000::2 source 1/1/2 vrf red
               switch(config)# no ipv6 route 6000::0/64 3000::3 source 1/1/2 vrf blue

      switch(config)# show runn
Current configuration:
!
      vrf blue
      vrf green
vrf red
      !

      vlan 1
                              interface 1/1/1
                              no shutdown
              vrf attach red
                              ip address 2000::1/64
                      interface 1/1/2
                      no shutdown
                       vrf attach green
                ip address 3000::1/64
                       interface 1/1/3
                      no shutdown
                      vrf attach blue
                      ip address 4000::1/64
```

Command History

Release	Modification
10.10	Inclusive language update.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config	Administrators or local user group members with execution rights for this command.

ip route interface

```
ip route <IPv4 ADDR/MEMBER> interface <IPv4 ADDR>
 [<DISTANCE> | <VRF INSTANCE-NAME>]
no ip route
```

Description

Creates a route leak between the SRC-VRF and DST-VRF. Using the static method, the route must first be added to the destination VRF. The route is added to the local interface of the source VRF with a next-hop interface. The existing IP route command can then take either the next-hop IP or the next-hop interface.

- **INTERFACE** refers to the outgoing interface in an m/s/p format
- Users must provide both the next-hop IP and the interface information to leak network routes.
- The next-hop IP information is not required to leak connected routes.

The `[no]` form of command deletes the static VRF leaked route.

Parameter	Description
<IPv4 ADDR/MEMBER>	Required: IPv4 IP-Address route destination.
<IPv4 ADDR>	Required: IPv4 route destination
<DISTANCE>	Optional: administrative distance of static route
<VRF INSTANCE-NAME>	Optional: VRF instance

Options

nullroute

Discard packets to the destined route silently.

reject

Discard packets to the destined route and return ICMP error to the sender.

Example

Configures a route leak between the SRC-VRF and DST-VRF:

```
switch(config)# show runn
Current configuration:
!
vrf blue
      vrf green
vrf red
      !

vlan 1
       interface 1/1/1
       no shutdown
              vrf attach red
       ip address 10.0.0.1/24
              interface 1/1/2
              no shutdown
              vrf attach green
       ip address 20.0.0.1/24
       interface 1/1/3
              no shutdown
       vrf attach blue
       ip address 40.0.0.1/24

      switch(config)# ip route A.B.C.D/M IPv4 route destination
      switch(config)# ip route A.B.C.D/M IPv4 route destination
      switch(config)# ip route 30.0.0.0/24 A.B.C.D Nexthop IPv4 address
      switch(config)# ip route 30.0.0.0/24 1/1/2 A.B.C.D Nexthop IPv4 address
```

```
      switch(config)# ip route 30.0.0.0/24 20.0.0.2 vrf green
      switch(config)# ip route 30.0.0.0/24 1/1/2 20.0.0.2 vrf red
      switch(config)# ip route 50.0.0.0/24 1/1/2 20.0.0.2 vrf blue
      switch(config)# ip route 50.0.0.0/24 20.0.0.2 vrf green
      switch(config)# ip route 60.0.0.0/24 1/1/2 vrf red

      switch(config)# show runn
Current configuration:
!
vrf blue
      vrf green
vrf red
      !

vlan 1
       interface 1/1/1
       no shutdown
              vrf attach red
       ip address 10.0.0.1/24
              interface 1/1/2
              no shutdown
              vrf attach green
       ip address 20.0.0.1/24
       interface 1/1/3
              no shutdown
       vrf attach blue
       ip address 40.0.0.1/24
       ip route 30.0.0.0/24 1/1/2 20.0.0.2 vrf red
       ip route 50.0.0.0/24 1/1/2 20.0.0.3 vrf blue
       ip route 60.0.0.0/24 1/1/2 vrf red


      switch(config)# no ip route 30.0.0.0/24 1/1/2 20.0.0.2 vrf red
      switch(config)# no ip route 50.0.0.0/24 1/1/2 20.0.0.3 vrf blue
      switch(config)# no ip route 60.0.0.0/24 1/1/2 vrf red

      switch(config)# show runn
Current configuration:
!
      vrf blue
      vrf green
vrf red
      !

      vlan 1
      interface 1/1/1
      no shutdown
      vrf attach red
      ip address 10.0.0.1/24
      interface 1/1/2
      no shutdown
      vrf attach green
      ip address 20.0.0.1/24
      interface 1/1/3
      no shutdown
      vrf attach blue
      ip address 40.0.0.1/24
```

Command History

Release	Modification
10.10	Inclusive language update.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

rd

```
rd <AS-NUMBER:NN>
no rd <AS-NUMBER:NN>
```

Description

Configures VRF table with specified route-distinguisher value. An RD ensures uniqueness of a route between multiple VRFs.

The **no** form of the command will delete RD from a specified VRF table. The VRF instance goes down when RD is deleted. All routes that are exported or leaked from the deleted VRF will be withdrawn.

Parameter	Description
<AS-NUMBER:NN>	Required: Enter an AS number and an arbitrary number.

Examples

Configures VRF for RD with an AS number 100:1.

```
switch(config-vrf)# rd 100:1
```

Deletes the RD from the specified VRF.

```
switch(config-vrf)# no rd
```

Deletes the RD and AS number from the specified VRF.

```
switch(config-vrf)# no rd 100:1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-vrf`	Administrators or local user group members with execution rights for this command.

redistribute

```
redistribute <protocol> [route-map <route-map-name>]
[no] redistribute <protocol> [route-map <route-map-name>]
```

Description

Specifies the protocol routes to redistribute to BGP VRF context. Any routes existing in the BGP VRF context are leaked as a VPNv4 or VPNv6 prefixes to other VRFs based on BGP route-targets.

The **no** form of this command removes the protocol.

Parameter	Description
`redistribute`	Required: redistributes routes from another routing protocol.
`connected`	Optional: redistribute directly attached networks.
`ospfv3`	Optional: redistributes OSPFv3 routes.
`static`	Optional: redistributes static routes.
`route-map`	Optional: applies route map policy for redistribution.

Examples

The following is an example of redistributing OSPFv2 routes to a BGP **`vrf cust_a`** instance by creating a router BGP instance for **cust_a**.

1. Creating the router BGP instance for **cust_a**.

```
switch(config)# router bgp 1
switch(config-router) # vrf cust_a
```

2. Redistributing the router to BGP.

```
switch(config-router-bgp)# redistribute ospf
```

The following is an example of redistributing OSPFv3 routes to a BGP `vrf cust_a` instance by creating a router BGP instance for `cust_a`.

1. Creating the router BGP instance for **`cust_a`**

```
switch(config)# router bgp 100
switch(config-router)# vrf cust_a
```

2. Configuring the address family IPv6 unicast to the router

```
switch(config-router-bgp)# address-family ipv6 unicast
```

3. Redistributing the router to OSPFv3

```
switch(config-router-ipv6-uc)# redistribute ospfv3
```

4. Redistributing the router configured with **`ipv6-af-us`** to OSPFv3

```
switch(config-router-bgp-vrf-ipv6-af-uc)# redistribute ospfv3
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

route-target

```
route-target [import | export | both] <AS-NUMBER:NN>
no route-target [import | export | both] <AS-NUMBER:NN>
```

Description

BGP route targets are extended BGP communities that identify the VPNv4 or VPNv6 routes that are associated with a VRF. This command specifies the route targets used on the import or export of the routes to other VRFs. Multiple route targets can be associated with a VRF.

The **no** form of the command removes the association.

Parameter	Description
`import`	Specifies the RTs imported to the VRF. Import or export or both required Literal Specifies the route-target type.
`export`	Specifies the RT on VPNv4 or VPNv6 prefixes that are leaked to other VRFs.
`both`	Specifies the RT for both export and import types.
<AS-NUMBER:NN>	Specifies an AS number and an arbitrary number for the RT value.

Examples

Configuring route targets for several VRFs.

```
switch(config)# vrf default
switch(config-vrf)# rd 192.168.2.1:0
switch(config-vrf)# address-family ipv4 unicast
switch(config-vrf-ipv4-af-uc)# route-target export 65001:0
switch(config-vrf-ipv4-af-uc)# route-target import 65001:1
switch(config-vrf-ipv4-af-uc)# route-target import 65001:2
switch(config-vrf-ipv4-af-uc)# exit-address-family
switch(config-vrf)# exit
switch(config)# vrf VRF1
switch(config-vrf)# rd 192.168.2.1:1
switch(config-vrf)# address-family ipv4 unicast
switch(config-vrf-ipv4-af-uc)# route-target export 65001:1
switch(config-vrf-ipv4-af-uc)# route-target import 65001:0
switch(config-vrf-ipv4-af-uc)# exit-address-family
switch(config-vrf)# exit
switch(config)# vrf VRF2
switch(config-vrf)# rd 192.168.2.1:2
switch(config-vrf)# address-family ipv4 unicast
switch(config-vrf-ipv4-af-uc)# route-target export 65001:2
switch(config-vrf-ipv4-af-uc)# route-target import 65001:0
switch(config-vrf-ipv4-af-uc)# exit-address-family
switch(config-vrf)# exit
```

Configuring the route target for export. Removing the configuration for export.

```
switch(config-vrf-ipv4-af-uc)# route-target export 100:1
switch(config-vrf-ipv4-af-uc)# no route-target export 100:1
```

Configuring the route target for import. Removing the configuration for import.

```
switch(config-vrf-ipv4-af-uc)# route-target import 100:2
switch(config-vrf-ipv4-af-uc)# no route-target import 100:2
```

Configuring the route target for both import and export. Removing the configuration for import and export.

```
switch(config-vrf-ipv4-af-uc)# route-target both 100:3
switch(config-vrf-ipv4-af-uc)# no route-target both 100:3
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Manager (#)	Administrators or local user group members with execution rights for this command.

router bgp

```
router bgp <AS-NUMBER>
no router bgp <AS-NUMBER>
```

Description

This command configures the BGP instance on the router, configures the AS (Autonomous System) the router belongs to, and enters into the BGP router configuration mode. Only a single BGP AS number can be assigned for the entire system.

The **no** form of the command deletes the BGP instance from the router.

Parameter	Description
AS-NUMBER	Specifies a 4-byte AS number in the range 1-4294967295 in integer format or from 0.1-65535.65535 in dotted format.

Examples

Configuring the BGP instance with the AS number:

```
```

Deleting BGP configurations:

```
switch(config)# no router bgp 100
This will delete all BGP configurations on this device.
Continue (y/n)?
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

route-target export

```
route-target export route-map <ROUTE-MAP-NAME>
no route-target export route-map <ROUTE-MAP-NAME>
```

Description

This command specifies the route-map to be used while exporting routes to VRF. Route target export route-map is configured to filter the routes those are leaked to VRF.

The **no** form of the command removes the specification.

Parameter	Description
`route-map`	Specifies route-map policy for export.
<ROUTE-MAP-NAME>	Specifies route-map policy name

VPNv4 SAFI is not supported.

Examples

Configuring route-map for route target export in IPv4:

```
switch(config)# vrf cust_a
switch(config-vrf)# address-family ipv4 unicast
switch(config-vrf-ipv4-af-uc)# route-target export route-map rmap
```

Removing the configuration for export in IPv4:

```
switch(config)# vrf cust_a
switch(config-vrf)# address-family ipv4 unicast
switch(config-vrf-ipv4-af-uc)# no route-target export route-map rmap
```

Configuring route-map for route target export in IPv6:

```
switch(config)# vrf cust_a
switch(config-vrf)# address-family ipv6 unicast
switch(config-vrf-ipv6-af-uc)# route-target export route-map rmap
```

Removing the configuration for export in IPv6:

```
switch(config)# vrf cust_a
switch(config-vrf)# address-family ipv6 unicast
switch(config-vrf-ipv6-af-uc)# no route-target export route-map rmap
```

Match prefix-list and set metric is only supported for route-target export route-map:

```
route-map rmap permit seq 1
       match ip address prefix-list prefx
       set metric 400
ip prefix-list seq 1 permit 14.1.1.0/24
```

Applies to both IPv4 and IPv6 address family only.

Only one route-target export route-map per VRF can be configured.

Command History

Release	Modification
10.12.1000	Command introduced.

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360	config-vrf	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
9300 9300S 10000		

router bgp vrf

```
router bgp <AS-NUMBER> [vrf <VRF-NAME>]
[no] router bgp <AS-NUMBER> [vrf <VRF-NAME>]
```

Description

This command configures VRF for the BGP instance.

The **no** form of this command removes the configuration.

Parameter	Description
AS-NUMBER	Specifies a 4-byte AS number in the range 1-4294967295 in integer format or from 0.1-65535.65535 in dotted format.
<VRF-NAME>	String VRF name for the VRF.

Usage

- Use the command **vrf *vrf-name*** within the router BGP context.
- **address-family {ipv4 | ipv6}** nodes are only supported within the VRF context.
- **address-family {ipv4 | ipv6}** nodes are required to redistribute the OSPF static/connected IPv4 or IPv6 routes.

Examples

Configure the VRF for customer A, on the BGP instance 100:

```
switch(config)# router bgp 100
switch(config)# vrf cust_a
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P	config	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8360 9300 9300S 10000		

show bgp vpn unicast

```
show bgp [vrf <VRF-NAME>][{ipv4 unicast | ipv6 unicast| ipv4 unicast| all unicast}] [vsx-peer]
```

Description

Shows the BGP-VPN per VRF routes with additional route information like RD and extended community route targets.

Displays the BGP neighbor information for the specified VRF.

By default the default_vrf BGP instance information is displayed if the VRF is not specified.

Parameter	Description
`unicast`	Selects the subaddress family identifier
`all`	Displays VPNv4 address family routes for all VRFs
`vrf`	Displays VPNv4 address-family routes for specified VRF
`vpn-addr-family`	Required: Literal Select the VPNv4 or VPNv6 address family
`vrf-name`	Required: Literal or string. Specify **all** or **vrf-name**.
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Examples

Required use of *`vpn-addr-family`*:

```
switch# show bgp <vpn-addr-family> unicast {all | [vrf <vrf-name> | A.B.C.D/M}
```

Show BGP VRF ipv4 unicast routes for vrf-name

```
switch# show bgp vrf Red ipv4 unicast
       Status codes: s suppressed, d damped, h history, * valid, > best, = multipath
       i internal, e external, S Stale, R Removed
       Origin codes: i - IGP, e - EGP, ? - incomplete

  VRF: Red
  Local Router-ID 172.16.3.1

  Network           Nexthop       Metric  LocPrf  Weight Path
```

```
  Route Distinguisher: 65000:1
  *> 172.16.0.0/24  0.0.0.0       0       100     32768      ?
  *> 172.16.1.0/24  0.0.0.0       0       100     32768      ?
  *> 172.16.2.0/24  172.16.0.2    0       100     32768      ?
  *> 172.16.3.0/24  172.16.0.3    0       100     32768      ?

  Total number of entries 4
```

```
switch# show bgp vrf Green ipv4 unicast
       Status codes: s suppressed, d damped, h history, * valid, > best, = multipath
       i internal, e external, S Stale, R Removed
       Origin codes: i - IGP, e - EGP, ? - incomplete

               VRF: Green
  Local Router-ID 172.17.2.1

  Network           Nexthop       Metric  LocPrf  Weight Path
  Route Distinguisher: 65000:2
  *> 172.17.0.0/24  0.0.0.0       0       100     32768      ?
  *> 172.17.1.0/30  0.0.0.0       0       100     32768      ?
  *> 172.17.2.0/24  172.17.0.2    0       100     32768      ?

  Total number of entries 3
```

```
switch# show bgp vrf Blue ipv4 unicast
       Status codes: s suppressed, d damped, h history, * valid, > best, = multipath
       i internal, e external, S Stale, R Removed
       Origin codes: i - IGP, e - EGP, ? - incomplete

  VRF: Blue
  Local Router-ID 172.18.3.1

               Network           Nexthop       Metric  LocPrf  Weight Path
  Route Distinguisher: 65000:3
  *> 172.18.0.0/24  0.0.0.0       0       100     32768      ?
  *> 172.18.1.0/30  0.0.0.0       0       100     32768      ?
  *> 172.18.3.0/24  172.18.0.3    0       100     32768      ?

  Total number of entries 3
```

```
switch# show bgp vrf Shared ipv4 unicast
       Status codes: s suppressed, d damped, h history, * valid, > best, = multipath
       i internal, e external, S Stale, R Removed
       Origin codes: i - IGP, e - EGP, ? - incomplete

  VRF: Shared
  Local Router-ID 192.168.99.1

               Network             Nexthop       Metric  LocPrf  Weight Path
  Route Distinguisher: 65000:99
  *> 192.168.99.0/24  0.0.0.0       0       100     32768      ?

  Total number of entries 1
```

Show BGP VRF ipv6 unicast routes for vrf-name:

```
switch# show bgp vrf Red ipv6 unicast
       Status codes: s suppressed, d damped, h history, * valid, > best, = multipath
       i internal, e external, S Stale, R Removed
       Origin codes: i - IGP, e - EGP, ? - incomplete

  VRF: Red
  Local Router-ID 172.16.3.1

  Network           Nexthop       Metric  LocPrf  Weight Path
  Route Distinguisher: 65000:1
  *  2001:100:1:1000:/56
        2001:100:1:1000::72a     0       0       200      ?
  *> 2001:100:1:1000::/56
           ::                    0       100     32768    ?
  *  2001:100:1:2000::/56
        ::FFFF:200.10.10.1 0     0       100     32768    ?

  Total number of entries 3
```

Show BGP VRF routes for all vrfs and all address-families:

```
switch# show bgp all-vrf all
       Status codes: s suppressed, d damped, h history, * valid, > best, = multipath
       i internal, e external, S Stale, R Removed
       Origin codes: i - IGP, e - EGP, ? - incomplete

  VRF: Blue
  Local Router-ID 172.18.3.1

  Address-family: IPv4 Unicast
                  ----------------------------
  Network           Nexthop       Metric  LocPrf  Weight Path
  *> 172.18.0.0/24  0.0.0.0       0       100     32768    ?
  *> 172.18.1.0/30  0.0.0.0       0       100     32768    ?
  *> 172.18.3.0/24  172.18.0.3    0       100     32768    ?

  Total number of entries 3

  Address-family: IPv6 Unicast
                  ----------------------------
  Network           Nexthop       Metric  LocPrf  Weight Path

  Toatl number of entries 0

  VRF: Green
  Local Router-ID 172.17.2.1

  Address-family: IPv4 Unicast
                  ----------------------------
  Network           Nexthop       Metric  LocPrf  Weight Path
  *> 172.17.0.0/24  0.0.0.0       0       100     32768    ?
  *> 172.17.1.0/30  0.0.0.0       0       100     32768    ?
  *> 172.17.2.0/24  172.17.0.2    0       100     32768    ?

  Total number of entries 3

  Address-family: IPv6 Unicast
                  ----------------------------
  Network           Nexthop       Metric  LocPrf  Weight Path
```

```
Total number of entries 0

VRF: Red
Local Router-ID 172.16.3.1

Address-family: IPv4 Unicast
             ----------------------------
Network           Nexthop       Metric  LocPrf  Weight Path
*> 172.16.0.0/24  0.0.0.0       0       100     32768    ?
*> 172.16.1.0/24  0.0.0.0       0       100     32768    ?
*> 172.16.2.0/24  172.16.0.2    0       100     32768    ?
*> 172.16.3.0/24  172.16.0.3    0       100     32768    ?

Total number of entries 4

Address-family: IPv6 Unicast
             ----------------------------
Network           Nexthop       Metric  LocPrf  Weight Path
*  2001:100:1:1000::/56
       2001:100:1:1000::72a     0       0       200      ?
*> 2001:100:1:1000::/56
         ::                     0       100     32768    ?
*  2001:100:1:2000::/56
       ::FFFF:200.10.10.1 0     0       100     32768    ?

Total number of entries 3

VRF: Shared
Local Router-ID 192.168.99.1

Address-family: IPv4 Unicast
             ----------------------------
Network             Nexthop       Metric  LocPrf  Weight Path
*> 192.168.99.0/24  0.0.0.0       0       100     32768    ?

Total number of entries 1

Address-family: IPv6 Unicast
             ----------------------------
Network           Nexthop       Metric  LocPrf  Weight Path

Total number of entries 0
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325	Operator (>)	

Platforms	Command context	Authority
8325H 8325P 8360 9300 9300S 10000		

show bgp info vrf

```
show bgp info vrf <vrf-name> [vsx-peer]
```

Description

Displays BGP route-targets information for specified VRF.

Parameter	Description
info	Display BGP RT information.
vrf-name	Required string VRF name for the vrf.
vsx-peer	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Example

Show BGP VRF information.

```
switch# show bgp info vrf red
 VRF : red
       VRF RD : 100:1

       Address-family IPv4 unicast info
       Redistribution : ospf
       Export RT list : 100:1 100:2
       Import RT list : 100:3

       Address-family IPv6 unicast info
       Redistribution : connected
 Export RT list : 100:11 100:12
 Import RT list : 100:15
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ip route vrf

```
show ip route vrf <vrf-name> [vsx-peer]
```

Description

Shows route information for specified VRF.

Parameter	Description
vrf-name	Required: string VRF name for the VRF.
vsx-peer	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Example

```
switch# show ip route vrf green

       Displaying ipv4 routes selected for forwarding
        '[x/y]' denotes [distance/metric]
               10.0.0.0/24, vrf green
                via 20.0.0.1[vrf red], [1/0], static
                30.0.0.0/24, vrf green
                via 1/1/2, [0/0], connected
   30.0.0.2/32, vrf green
                via 1/1/2, [0/0], local
                60.0.0.0/24, vrf green
                via 1/1/1[vrf red], [1/0], static
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ipv6 route vrf

```
show ipv6 route vrf <VRF-NAME> [vsx-peer]
```

Description

Shows the route information for specified VRF.

Parameter	Description
<VRF-NAME>	Required: String VRF name for the VRF.
vsx-peer	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Examples

Displaying ipv4 routes selected for forwarding

```
switch(config)# show ip route vrf green
Origin Codes: C - connected, S - static, L - local
R - RIP, B - BGP, O - OSPF, D - DHCP
U - Unnumbered
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
IA - OSPF internal area, E1 - OSPF external type 1
E2 - OSPF external type 2
VRF: green
Prefix              Nexthop         Interface     VRF(egress)       Origin/
Distance/    Age
Type      Metric
------------------------------------------------------------------------------
20.1.1.0/24         30.1.1.1        1/1/2         red               S
[1/0]        00h:00m:04s
30.1.1.0/24         -               1/1/2         red               B/I
[200/0]      00h:01m:05s
40.1.1.0/24         -               1/1/3         -                 C
[0/0]        -
40.1.1.1/32         -               1/1/3         -                 L
[0/0]        -
Total Route Count : 4
switch(config)#
switch(config)#  show ip route vrf red
Displaying ipv4 routes selected for forwarding
```

```
Origin Codes: C - connected, S - static, L - local
R - RIP, B - BGP, O - OSPF, D - DHCP
U - Unnumbered
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
IA - OSPF internal area, E1 - OSPF external type 1
E2 - OSPF external type 2
VRF: red
Prefix              Nexthop      Interface     VRF(egress)       Origin/
Distance/    Age
Type      Metric
------------------------------------------------------------------------------
30.1.1.0/24         -            1/1/2         -                 C         [0/0]
     -
30.1.1.1/32         -            1/1/2         -                 L         [0/0]
     -
40.1.1.0/24         -            1/1/3         green             B/I       [200/0]
     00h:02m:42s
Total Route Count : 3
```

Displaying ipv6 routes selected for forwarding

```
Switch# show ipv6 route all-vrfs
Origin Codes: C - connected, S - static, L - local
R - RIP, B - BGP, O - OSPF, D - DHCP
U - Unnumbered
Type Codes:   E - External BGP, I - Internal BGP, V - VPN, EV - EVPN
IA - OSPF internal area, E1 - OSPF external type 1
E2 - OSPF external type 2
VRF: green
Prefix                          Nexthop     Interface     VRF(egress)       Origin/
  Distance/    Age
Type      Metric
------------------------------------------------------------------------------------
----------------------------
2000::/64                       3000::1     1/1/2         red               S
  [1/0]        00h:01m:05s
3000::/64                       -           1/1/2         red               B/I
  [200/0]      00h:00m:16s
4000::/64                       -           1/1/3         -                 C
  [0/0]        -
4000::1/128                     -           1/1/3         -                 L
  [0/0]        -
VRF: red
Prefix                          Nexthop     Interface     VRF(egress)       Origin/
  Distance/    Age
Type      Metric
------------------------------------------------------------------------------------
----------------------------
3000::/64                       -           1/1/2         -                 C
  [0/0]        -
3000::1/128                     -           1/1/2         -                 L
  [0/0]        -
4000::/64                       -           1/1/3         green             B/I
  [200/0]      00h:00m:11s
Total Route Count : 7
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

vrf

```
vrf <VRF-NAME>
no vrf <VRF-NAME>
```

Description

Creates a VRF instance named <*VRF-NAME*> and then enters its context. Use **default** for <*VRF-NAME*> to enter the default VRF configure context.

Except for the default VRF, the **no** form of the command deletes the named VRF instance and any IP configuration for interfaces or SVI linked to default VRF. The default VRF cannot be deleted and a warning is given if attempted. To erase the Route-Distinguisher and Route-Targets, enter the default VRF context and delete them manually one by one.

Only default and mangement VRFs are supported on the following 6300 switches: S3L75A, S3L76A, S3L77A.

Parameter	Description
<*VRF-NAME*>	Specifies the VRF name. Range: Up to 32 alphanumeric characters. The **mgmt** VRF cannot be used.

Examples

Creating the VRF named **cust_A** and then entering its context:

```
switch(config)# vrf cust_A
```

Entering the **default** VRF context:

```
switch(config)# vrf default
```

Deleting the VRF named **test**:

```
switch(config)# no vrf test
```

Command History

Release	Modification
10.09	Added default VRF information.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config	Administrators or local user group members with execution rights for this command.

Chapter 12
Policy Based Routing (PBR)

Policy Based routing (PBR) lets you manipulate the path of a packet based on the various attributes of the packet. Packets that can be manipulated by PBR are packets that are already routing through the system at Layer 3, with a destination IP address that is on a network other than the packet ingress Layer 3 interface. PBR is an extension of the existing classifier policy system where traffic to be manipulated is matched by a classifier **class**, and policy **actions** to be executed on the matching traffic. Matching traffic with the same destination can be routed over different paths so that different types of traffic, such as VoIP or traffic with special security requirements, can be better managed.

PBR's ability to influence a packet path is limited to the current router in that the next router to which the traffic is redirected makes an independent decision on where to forward traffic to next, if anywhere. Packets not matched by a PBR policy entry are not affected and take routes specified in the system route table.

PBR actions

- **interface null:** equivalent to the policy **drop** policing action. Any packets matching the class criteria for that policy entry will be dropped and not routed any further.
- **interface tunnel:** allows for specifying a GRE, 6in4 or 6in6 tunnel as the outbound interface for all matching packets. The tunnel must exist before configuring. Packets sent into the tunnel interface egress at the router at the endpoint of the tunnel. If the tunnel is misconfigured or down the traffic may be lost.
- **nexthop:** allows for overriding the routing table's longest prefix match next-hop router for matching packets. If no such routing table entry exists for matching packets, (default or not) this action still affects matching packets.
- **default-nexthop:** allows for specifying a next-hop router for matching packets when there is no longest prefix match for those packets in the routing table. Such a default-nexthop overrides a system default route if already configured and also applies if there is no system default route.

- Next-hop and default-nexthop facilitate routing matching packets where they otherwise might not be, due to the absence of routing table entries.
- Unlike next-hop, default-nexthop only applies if there is no destination lookup match in the main routing table for matching packets.

PBR policy action and action list

Entries in a policy use a class to specify the criteria on what packets to match and a set of one or more policy actions to take on the matching packets such as dscp or mirror.

PBR is configured for use as another such policy action in a policy entry by using the `pbr` keyword followed by the name of a user configured PBR action list. This list may contain up to eight PBR action

entries which are listed according to priority by sequence number. The PBR actions section lists the different action entry choices.

In a given action list applied to an interface, zero entries may be available (that is, reachable). Of the entries that are available, only one can be active for that action list in a given VRF. Due to real-time changes in the network operating environment, list entry availability can change at any moment. At any given time, the highest priority available entry in the list is selected as the active PBR action for that list. If the active entry becomes unavailable, the next highest available entry is automatically selected and replaces that now-unavailable entry as the active entry. If a higher priority entry (other than the current active entry) becomes available, this entry is selected and replaces the current active entry, as the new active entry (even though that current entry is still available). If no entries are available or the only available entry becomes unavailable, routing occurs based on the routing table and is not influenced by PBR.

The **interface null** action is always considered available. If there are no higher priority PBR actions available, it will be selected as the active entry when configured in a PBR action list, and applied to an interface.

PBR action list maximum entries

There is a limit of eight entries per action list and a limit of 512 unique entries across the entire system. This value restricts the sum of all unique next-hops, default-nexthops, tunnel interfaces, and null interfaces applied to all Layer 3 interfaces across all VRFs. A given set of entries that are unique in one VRF count against this limit if the set is applied in another VRF.

For example, a given next-hop specified in an action list applied to an interface counts against the limit once. If that same next-hop address is entered in a different action list and applied to another interface in the same VRF, this next-hop is not unique in the VRF and does not affect the limit further. If one of those action lists (that specify the address) is applied to an interface in a different VRF, the entry will affect the limit since addresses are unique to each VRF.

Another example is if two PBR action lists with eight unique next-hop entries (for a total of 16 unique next-hops) are configured as PBR actions in a policy. If that policy is applied to 32 route-only ports (each of which is in a different VRF), the entry supply will be exhausted.

- The specified capabilities and capacities limit for a unique VRF is 33.

IP versions in an action list

Configuring PBR actions of different IP versions in the same PBR action list is not supported. This limitation applies to the next-hop and default-nexthop actions specifically. If an action of one IP version is already configured, configuring an action with a different IP version in the same action list will be blocked. This limitation applies regardless of whether the action list is part of a policy, applied to an interface or whether any of the entries are selected or not.

Furthermore, use cases similar to the following are not supported:

- Applying a PBR action list with entries of one IP version (for example IPv4 next-hop) to an interface of the other IP version (for example a route-only port with an IPv6 address only).
- Creating a policy entry with a class of one IP version and a PBR action list with entries of the other IP version.

Configuration mismatches while not prevented are not supported and behavior is not defined. At best, traffic will not reach its intended destination.

Specifying valid next-hop and default-nexthop addresses

PBR supports all types of directly connected hosts as nexthop or default-nexthop actions. PBR also supports recursive nexthops, however only for remote hosts that are learned over a VXLAN L3VNI tunnel. Refer toPBR and VXLAN for more details.

This means that while you may specify any IPv4 or IPv6 address as a nexthop or default-nexthop, afrer application only those reachable through a directly connected network will be installed as an active PBR nexthop. Even if a specified address is reachable from the router, if it is not on a directly connected subnet it will not be installed. This is also true for nexthops reachable through a VXLAN L3VNI tunnel.

Hardware path PBR versus software path PBR

Traffic to be routed that cannot be handled by the switching ASIC (for example IPv4 packets with IP options), is handled by the switch operating system kernel routing software.

The characteristics of network usage determine the proportion of traffic handled by software path PBR relative to the hardware path. Under normal conditions, software path PBR is likely to range from a fraction of a percent of all traffic to almost none.

Packets with IP options are sent to the hardware ASIC as well as the CPU. Since the software path PBR actions are applied to the packets in the CPU, a default entry is programmed in the hardware to take no action on those packets. This hardware entry count can be seen in diagnostics with the command `diag-dump policy basic`.

Hardware versus software path for default-nexthop action

When a policy containing an entry with PBR is applied to an interface and the active action for that entry action list is 'default-nexthop', packets that match the class criteria and have no explicit route table hit (that is, no destination address longest prefix match, a.k.a 'route-miss'), they are forwarded to the specified active PBR default-nexthop. This is true for packets that follow the hardware and software paths through the switch.

There is a difference between PBR hardware and software path behavior when a route table hit occurs for class-matched packets (when the policy entry is a PBR default-nexthop).

Hardware path match criteria for a PBR policy entry with default-nexthop is extended to include the route table miss along with the class qualifiers, resulting in a **policy entry hit** for that policy entry. Conversely when there is a route table hit, the result is a policy entry miss in hardware path. When a policy entry miss occurs, policy processing moves on to the next entry in the policy and takes whatever action is specified, if any exist. This can include a different PBR routing action including interface null or no PBR action at all, for example.

In software path, the class match criteria is the only criteria required to achieve a policy entry hit. When that occurs, policy processing will stop. When there is a class match and a route table hit (with PBR action default-nexthop), the packet is forwarded according to that route table entry, not the PBR default-nexthop entry, nor is it influenced by any subsequent policy entries.

This difference in behavior is due to a limitation in the software path matching and routing engine, relative to hardware.

Table 1: *default-nexthop behaviors*

Routing packet	System route miss	System route hit
Hardware	Take PBR default-nexthop route.	Policy entry miss, continue policy processing with next entry, if present. No further matches result in system default route being used, if present.
Software	Take PBR default-nexthop route.	Policy entry hit is overridden by system route hit. System route entry used, policy processing stops.

Software path and system default route

AOS-CX default behavior (for packets with destination networks that lack entries in the system route table with a reachable next-hop address), sends packets to the system CPU for special handling, if any are specified (for example, sending an ICMP unreachable reply message).

A side effect of this CPU routing is that it has higher priority than an applied PBR next-hop action. Even though a policy is applied with an entry that matches the traffic and specifies a PBR next-hop action which is reachable, the traffic will still be routed to the system CPU (due to the absence of a reachable next-hop in the system route table), and not through the desired PBR hardware path. With traffic routing to the system CPU, it will be properly routed by PBR software path, but it will also be rate limited by the control plane policing feature. Loss can occur at higher traffic rates.

The workaround for this issue is to create a default next-hop route in the system with a reachable next-hop router/host. This will result in a route hit, or reachable next-hop detected for the traffic with no further need to route traffic to the CPU. The PBR hardware path next-hop action will then occur, as desired.

PBR and VRFs

A given physical router can be partitioned into multiple virtual routers. All frontplane ports and logical interfaces are initially in the default VRF. Any subset of frontplane ports or logical interfaces can be selectively moved into a user configured VRF, to establish a routing topology that is different from the default VRF.

Policies with PBR actions can be applied to interfaces in any VRF, and PBR action lists can be used in different policies across different VRFs. The effect of applying the same Policies/PBR action lists across different VRFs depends on the IP networks and interfaces configured in the different VRFs. For example, an action list that specifies actions of 'nexthop 1.1.1.10' and 'interface tunnel gre_10' could be used as the PBR action parameter for an entry in policy_1 and also in policy_2. If policy_1 is applied to an interface in VRF 'red', which has an interface in subnet 1.1.1.0/24 but no GRE tunnel named 'gre_10', then only the 'nexthop' action will be relevant to VRF 'red'. If policy_2 (which contains the same action list), is applied to an interface in VRF 'blue' (which lacks the 1.1.1.0/24 subnet configured but does have a tunnel named 'gre_10'), then only that interface tunnel action will apply in that VRF.

- It is possible to configure the same subnet in different VRFs, however named tunnel interfaces can only exist in one, so in the example of a common action list, the 'next-hop' action could be relevant to both VRFs, but the 'interface tunnel' action may only be relevant to one. If VRFs are part of the router configuration, be mindful of them when creating and applying policies with PBR action lists and their entries.
- VRF Route Leaking is not supported in the current release of PBR (10.4).
- Show commands can only reference default VRF.

PBR, ECMP, and routing protocols

When a policy with PBR actions is applied to an interface, the highest priority action from the action list is applied to matching traffic on that interface and overrides any static, ECMP or dynamically installed routes in the system.

PBR, VSX, and VLAN ACLs

Multichassis Link Aggregation Group (MCLAG) with VSX is a high-availability feature where a switch containing LAG members is connected to multiple switches to allow for node-level redundancy on that link. If one of the other switches goes down the LAG remains up and can continue to carry all the LAG traffic, bandwidth permitting.

A LAG (and an MCLAG) can be a member of a VLAN and a Layer 3 VSI (virtual switch interface), which can be created for that VLAN for the purposes of routing a policy with Layer 3 specific actions (that is, PBR), which can therefore be applied to that interface to influence routing decisions for matching Layer 3 packets on the MCLAG; like any other route-only port or VSI. Under such a configuration, it is likely that it is desirable to use VLAN ACLs applied to the VLAN the MCLAG is a member of.

There is a limit with this particular combination when the VLAN ACL specifies both IPv4 and IPv6 entries and the PBR policy has entries with IPv4 and IPv6 classes. This configuration will exhaust the switching ASIC resources and will fail to apply. To achieve a similar configuration, you may use port ACLs for IPv4 traffic (applying the ACL to all ports individually in the VLAN) while preserving the rest of the configuration.

PBR software path, VSX, and VRRP

Due to the dynamic nature of the VSX and VRRP protocols, applying a policy with PBR to VSX or VRRP Layer 3 kernel interfaces and then making further changes to those protocols may not result in expected behavior. Apply policies with PBR to VSX and VRRP Layer 3 interfaces after all changes have been made. If further changes are necessary the policy should be removed first and reapplied when configuration updates are complete.

PBR and next-hop router reachability

When a policy with PBR actions is applied to an interface, all nexthop and default-nexthop entries in the associated action list are continuously monitored for reachability. A probe occurs every 5 seconds for each entry. These probes guarantee that the reaction to the loss of reachability or detection of the new reachability of any such entry will be 5 seconds or less.

For example, if there are two reachable nexthops in an applied action list and the active nexthop (highest priority by lowest sequence number) entry becomes unreachable, the switchover to the lower priority nexthop will occur within 5 seconds. If there is only one nexthop or default-nexthop entry and no entries of any other type in the action list, routing decisions will return to the system routing table within 5 seconds.

Conversely if an applied action list has a single unreachable nexthop and that nexthop at some point becomes reachable, the switch back to routing to that nexthop will occur within 5 seconds of the new reachability of the nexthop on the network. Or in another scenario an action list has multiple nexthops and the active nexthop is not the highest priority entry. When the nexthop with the higher priority becomes available, the promotion of that now reachable nexthop to active will occur within 5 seconds of it's new reachability on the network.

One caveat is if the nexthop is a remote host learned over a VXLAN L3VNI tunnel. This is the only scenario in which a reursive nexthop is able to be programmed. In this scenario, the periodical probing of this nexthop is not performed and so there is no guarantee of reachability. The connected nexthop to reach the recursive nexthop will be probed. The recursive nexthop will be used as long as there is a EVPN-BGP route for the nexthop's subnet. For more information on this specific scenario, refer to the PBR nexthop over VXLAN L3VNI.

PBR and VXLAN

PBR supports configuring remote hosts as nexthops that are learned over a Virtual eXtensible LAN (VXLAN) Tunnel. When PBR is configured on a VTEP device, irrespective of the VXLAN configuration (Static VXLAN or VXLAN with BGP EVPN), remote hosts whose MAC/ARP entries are learned by the VTEP device can be configured as nexthops.

PBR nexthop over VXLAN L3VNI

In general PBR does not support remote (recursive) nexthops. However, PBR supports configuring remote non-connected hosts as PBR nexthops or default-nexthop actions that are learned over EVPN-distributed subnets across a routed VXLAN L3VNI tunnel. These nexthops are not guaranteed to exist, but are supported as a PBR nexthop or default-nexthop action as long as there is a resolved EVPN-BGP route that matches the PBR nexthop address. As a result, hosts learned over EVPN-distributed subnets across a routed VXLAN L3VNI tunnel will not be probed for reachability.

For more information on VXLAN or EVPN, refer to the VXLAN Guide.

A use-case of PBR over VxLAN L3VNI is to route packets to a dedicated machine over a tunnel such as a firewall or load balancer. The incoming packet will hit a VTEP and PBR policy will send the packet into the tunnel (ingress). On the other end, the packet will exit the tunnel and PBR policy will route the packet to the correct nexthop (egress). It is important to note that

In order to configure the end-to-end solution, a policy needs to be configured on both VTEPs, ingress and egress.

On 6300, 6400, and 8360 based platforms with PBR nexthop over VxLAN L3VNI configured, routing packets via software path is not supported. Packets that are not hardware-routed (such as packets that are forced to be software routed due to IP options) will be dropped. This is due to a limitation on the ASIC in these platforms.

Refer to Hardware versus software path for default-nexthop action for an explanation on the difference between software path and hardware path.

Ingress policy

In order to route traffic to a nexthop learned over VXLAN L3VNI, a policy needs to be applied to the ingress VTEP's L3 interface associated with the arriving traffic that needs to be redirected.

In the following example the policy is applied to **vlan 20**, (where **nexthop 100.1.1.4** is learned over VxLAN L3VNI):

```
class ip class1
    10 match any any any
pbr-action-list pbr1
    10 nexthop 100.1.1.4
policy p1
    10 class ip class1 action pbr pbr1
interface vlan 20
    vrf attach red
    ip address 200.1.1.2/16
    active-gateway ip mac aa:ee:ff:00:00:20
    active-gateway ip 200.1.1.1
    apply policy p1 routed-in
```

Egress policy

A policy must be applied to the egress VTEP's L3VNI that is receiving the traffic from the VXLAN tunnel.

In the following example the policy is applied to VNI 10000:

```
class ip class1
    10 match any any any
pbr-action-list pbr1
    10 nexthop 100.1.1.4
policy p1
    10 class ip class1 action pbr pbr1
interface vxlan 1
    source ip 1.1.1.1
    no shutdown
    vni 10
        vlan 10
    vni 10000
        vrf red
        routing
        apply policy p1 routed-in
```

L2VNI vs L3VNI

There is an important distinction between PBR nexthops learned over L2VNI and L3VNI. Any L2VNI is treated the same as a locally connected nexthop. It will be probed periodically for reachability. If this L2VNI nexthop is determined to be unreachable, then PBR will take action on the next item in the action list. If the nexthop becomes reachable again and it has a higher precedence than the currently configured action, then it will take over. See PBR policy action and action list for more details.

For recursive nexthops (learned over L3VNI), the connected nexthop to reach the recursive nexthop is constantly probed for reachability. However, the recursive PBR nexthop itself will not be periodically probed. They are not guaranteed to be reachable. The PBR action list will use the nexthop as long as there is a EVPN-BGP route that matches.

PBR and subinterfaces

Certain types of Layer 3 physical interfaces (ROP, L3 LAG, Split interface) can be divided into multiple logical Layer 3 interfaces for sending and receiving data, tagged with different 802.1Q VLAN IDs.

Policies with PBR actions cannot be applied to the parent physical interface of subinterfaces but may be applied to the subinterfaces themselves.

It is not necessary to configure VLANs on the switch before specifying them as subinterface encapsulation IDs.

Example

Creating a subinterface on a route-only port specifying an 802.1Q VLAN encapsulation ID and applying a policy with PBR action to the subinterface:

```
switch(config)# interface 1/1/1
switch(config-if)# routing
switch(config-if)# no shutdown
switch(config-if)# exit

switch(config)# interface 1/1/1.100
switch(config-subif)# encapsulation dot1Q 100
switch(config-subif)# ip address 100.1.1.1/24
switch(config-subif)# apply policy p1 routed-in
switch(config-subif)# exit
```

CLI errors

Table 1: *Error messages and triggering events*

Message	Event
`Failed to create action list`	A PBR action list create operation was unable to create the OVSDB row.
`PBR action list '%s' doesn't exist`	A PBR action list update or delete operation was unable to locate the list OVSDB row. OR A PBR action list entry create, update, or delete operation was unable to locate its parent action list OVSDB row.
`Unable to automatically set sequence number`	The automatically generated sequence number for a new PBR action list entry is higher than the maximum permitted sequence number.
`Unable to add PBR action list entry`	A PBR action list entry create failed to create an entry row in OVSDB.
`Configuration of IPv4 and IPv6 addresses in the same PBR action list is not supported`	An attempt to create both IPv4 and IPv6 addresses in the same action list was made.

Message	Event
`Invalid sequence number`	A PBR action list entry delete did not specify the entry sequence number.
`Cannot add entry because its automatic sequence number would exceed the maximum`	The automatic sequence number exceeded the maximum.
`Action list entry does not exist`	A PBR action list entry delete could not locate the OVSDB entry row.
`Action list name required`	A PBR action list delete was passed an invalid list name.
`Invalid action list name length`	A PBR action list name was longer than the permitted length.

Backup nexthop groups

In a network with thousands of routers exchanging millions of routes, many routes are reachable via more than one nexthop. It is not uncommon to see multiple routes reachable via the same list of nexthops. Given the large scale of routes, it is desirable to restore traffic after a failure in a time period that does not depend on the number of routes.

BGP routing protocol supports Prefix Independent Convergence (PIC). BGP PIC creates and stores a backup or secondary path in the routing information base (RIB). When a failure is detected, the backup path immediately takes over enabling fast convergence.

- Primary and secondary nexthop pair are referred to as *backup nexthop group*.
- By defining a *static backup nexthop* group, the user is not dependent on protocol based PIC and instead is relying on manual configuration to achieve faster convergence for a large number of routes.

PBR commands

apply policy

```
apply policy <POLICY-NAME> routed-in
no apply policy <POLICY-NAME> routed-in
```

Description

Applies a classifier policy containing a PBR action to an interface. A policy with PBR actions is only applicable to L3/routing interfaces.

The **no** form of this command removes a classifier policy containing a PBR action from an interface.
`config-if`

Parameter	Description
<POLICY-NAME>	Specifies name of the policy.

Restrictions

- Only Layer 3 interfaces are valid for PBR policy application, and only in the routed inbound direction.
- If a policy with an 'interface tunnel' PBR action is applied on a Layer 3 interface in VRF 'A', and that interface tunnel is a member of VRF 'B', the interface tunnel is considered down/unavailable in this policy application in VRF 'A'.

Usage

To use route-only ports (ROPs) as Layer 3 interfaces, an internal VLAN range must be configured first. A policy with PBR actions can be applied to ROPs.

Example

On the HPE Aruba Networking 6400 Switch Series, interface identification differs.

Applying a policy to an interface:

```
switch(config)# interface 1/1/10.10
switch(config-if)# routing
switch(config-if)# apply policy pbr_policy routed-in
switch(config-if)# exit
```

Applying a policy to a subinterface, inbound direction:

```
switch(config)# interface 1/1/10.10
switch(config-if)# apply policy my_policy in
switch(config-if)# exit
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ip nexthop group type backup

```
ip nexthop group <GROUP-NAME> type backup
no ip nexthop group <GROUP-NAME> type backup
```

Description

Creates backup nexthop group with a given name. If the backup nexthop group already exists, this command enters that backup nexthop group's configuration context.

The **no** form of this command deletes the backup nexthop group context and all configuration related to the backup nexthop group.

Parameter	Description
<GROUP-NAME>	Specifies nexthop group name. Nexthop group name can be a maximum of 64 alphanumeric characters.

Example

Creating a backup nexthop group named **grp1**:

```
switch(config)# ip nexthop group grp1 type backup
switch(config-nexthop-group-grp1)#
```

Deleting a backup nexthop group named **grp1**:

```
switch(config)# ip nexthop group grp1 type backup
switch(config-nexthop-group-grp1)# no ip nexthop group grp1 type backup
```

Command History

Release	Modification
10.11	Command introduced

Command Information

Platforms	Command context	Authority
9300 9300S	config	Administrators or local user group members with execution rights for this command.

pbr-action-list

```
pbr-action-list <ACTION-LIST-NAME>

    [<SEQUENCE-NUMBER>]
      {nexthop | default-nexthop} <NEXT-HOP-IP-ADDR>
       interface {null | <TUNNEL-NAME>}

    no [<SEQUENCE-NUMBER>]
      {nexthop | default-nexthop} <IP-ADDR>
       interface {null | <TUNNEL-NAME>}

no pbr-action-list <ACTION-LIST-NAME>
```

Description

Creates a PBR action list or modifies its entries.

The **no** form of this command can be used to delete an action list or an individual action list entry.

Parameter	Description
<ACTION-LIST-NAME>	Specifies the action list name. An action list name can be 1 to 64 alphanumeric characters.
<SEQUENCE-NUMBER>	Specifies list entry sequence number. Range: 1-4294967295 {nexthop | default-nexthop} Selects a regular next-hop (**nexthop**) or a default next-hop (**default-nexthop**). These parameters specify the address of a next-hop router to forward traffic matched by a class under different conditions.
`nexthop`	Sets the next hop for routing the packet.
`default-nexthop`	Sets the next hop for routing the packet when there is no explicit route for its destination.
<NEXTHOP-IP-ADDR>	Specifies IPv4 or IPv6 address of the next-hop router.
`interface {null |` *<TUNNEL-NAME>*	Selects the type of keyword interface: **null** or the tunnel interface name.
`interface null`	Selects keyword interface: **null**.
`null`	Specifies to drop matching traffic.
<TUNNEL-NAME>	Specifies an IP tunnel interface name through which to forward the matching traffic.

Restrictions

The reachability of the next-hop routers/tunnel interfaces in the list is not guaranteed. Such reachability can change at any time due to the dynamic nature of the network environment.

Usage

Each action list may contain up to eight entries of four different entry types:

- `interface null`
- `interface tunnel`
- `nexthop`
- `default-nexthop`

List entries have a unique sequence number which, if not user specified, are automatically assigned beginning at 10 and continuing at intervals of 10 for each subsequent new list entry, for example 20, 30, and 40. Sequence numbers of any value can be specified manually, a different interval may be set, and new entries can be added to (or removed from) any location in the list at any time.

Specifying an existing sequence number causes the existing list entry to be replaced by the new details. The list entry with the lowest sequence number has the highest priority entry in the list. The sequence numbers may be renumbered with the pbr-action-list resequence command.

Only one next-hop router or interface from the list is used per packet matched. This router or interface is defined as the highest priority list entry that is reachable or available at the time of the traffic match. If the highest priority list entry next-hop router or tunnel interface is reachable - that list entry is chosen, the search is stopped, and the traffic is forwarded to the next-hop router or interface for the entry. If the highest priority list entry next-hop router or tunnel interface is not reachable, the next highest priority list entry reachability is determined and used if reachable, otherwise the process continues

down the list. If none of the routers in the list are reachable, the packet may be dropped (through the null interface entry if configured) or forwarded according to a system route table entry.

An action list that contains a next-hop of one IP version cannot also contain an entry of another IP version. For example, an action list must contain only IPv4 or IPv6 next-hop addresses or tunnel interfaces.

Examples

The list name is included in the context prompt for easy current-list identification. Any list name over 10 characters will be truncated at 10 characters and terminated with the tilde character (~) to indicate a reduced list name display. This reduction affects the prompt display of the list name only:

```
switch(config)# pbr-action-list eighteenchars
switch(config-pbr-action-list-eighteench~)#
```

The following example creates an action list with two IPv4 next-hops, a default IPv4 next-hop, and a null interface. The example uses default sequence numbering for its list entries.

```
switch(config)# pbr-action-list test1
switch(config-pbr-action-list-test1)# nexthop 1.1.1.1
switch(config-pbr-action-list-test1)# nexthop 2.2.2.2
switch(config-pbr-action-list-test1)# default-nexthop 9.9.9.9
switch(config-pbr-action-list-test1)# interface null
switch(config-pbr-action-list-test1)# end

switch(config)# show pbr-action-list test1


           Name
  Sequence Type                              Address/Interface
-------------------------------------------------------------------
           test1
       10  nexthop                           1.1.1.1
       20  nexthop                           2.2.2.2
       30  default-nexthop                   9.9.9.9
       40  interface                         null
```

The following example creates an action list with an IPv4 next-hop and a tunnel interface with manual sequence numbers for its entries.

```
switch(config)# pbr-action-list test2
switch(config-pbr-action-list-test2)# 6 ip default-nexthop 4.4.4.4
switch(config-pbr-action-list-test2)# 1 interface tunnel10
switch(config-pbr-action-list-test2)# end

switch(config)# show pbr-action-list test2

           Name
  Sequence Type                              Address/Interface
-------------------------------------------------------------------
           test2
        1  interface                         tunnel10
        6  default-nexthop                   4.4.4.4
```

The following example creates an action list with two IPv4 tunnel interfaces, with default sequence numbering.

```
switch(config)# pbr-action-list test3
switch(config-pbr-action-list-test3)# interface tunnel10
switch(config-pbr-action-list-test3)# interface tunnel15
switch(config-pbr-action-list-test3)# end

switch(config)# show pbr-action-list test3

           Name
  Sequence Type                            Address/Interface
-------------------------------------------------------------------
           test3
       10  interface                        tunnel10
       20  interface                        tunnel15
```

The following example creates an action list with two IPv6 next-hops and the null interface, with manual sequence numbers.

```
switch(config)# pbr-action-list test4
switch(config-pbr-action-list-test4)# 5 nexthop 2000:abcd::cccc:dddd
switch(config-pbr-action-list-test4)# 6 nexthop 1000:abcd::1234:5678
switch(config-pbr-action-list-test4)# 7 interface null
switch(config-pbr-action-list-test4)# end

switch(config)# show pbr-action-list test4

           Name
  Sequence Type                            Address/Interface
-------------------------------------------------------------------
           test4
        5  nexthop                          2000:abcd::cccc:dddd
        6  nexthop                          1000:abcd::1234:5678
        7  interface                        null
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config` The `pbr-action-list` *<ACTION-LIST-NAME>* command takes you into the `config-pbr-action-list-` *<ACTION-LIST-NAME>* context where you modify entries for a PBR action list.	Administrators or local user group members with execution rights for this command.

pbr-action-list copy

`pbr-action-list <ACTION-LIST-NAME> copy <DESTINATION-ACTION-LIST-NAME>`

Description

Copies an existing PBR action list.

Parameter	Description
<ACTION-LIST-NAME>	Specifies the action list name to be copied.
<DESTINATION-ACTION-LIST-NAME>	Specifies the name of the copied action list. A destination action list name can be 1 to 64 alphanumeric characters.

Examples

The following example copies test4 action list to test 5.

```
switch(config)# show pbr-action-list test4

           Name
  Sequence Type                              Address/Interface
-------------------------------------------------------------------
           test4
         5 nexthop                            2000:abcd::cccc:dddd
         6 nexthop                            1000:abcd::1234:5678
         7 interface                          null

switch(config)# pbr-action-list test4 copy test5
switch(config-pbr-action-list-test4)# show pbr-action-list test5

           Name
  Sequence Type                              Address/Interface
-------------------------------------------------------------------
           test4
         1 nexthop                            2000.abcd::cccc.dddd
        11 nexthop                            1000.abcd::1234.5678
        21 interface                          null
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325	`config`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325H 8325P 8360 9300 9300S 10000		

pbr-action-list resequence

`pbr-action-list <ACTION-LIST-NAME> resequence <STARTING-SEQUENCE-NUMBER> <INCREMENT>`

Description

Renumbers the entries in an action list. The list entry with the lowest sequence number has the highest priority entry in the list.

Parameter	Description
<ACTION-LIST-NAME>	Specifies the action list name to have its entries resequenced.
<STARTING-SEQUENCE-NUMBER>	Specifies the starting sequence number. Range: 1-4294967295
<INCREMENT>	Specifies the increment of the resequencing. Range: 1-4294967295

Examples

The following command shows how a PBR action list is resequenced. In the following example, an action list named **test4** is resequenced so that instead of its entries starting at 5 and being numbered sequentially, its entries start now at 1 and they are numbered in increments of 10:

```
switch(config)# show pbr-action-list test4

           Name
  Sequence Type                              Address/Interface
-------------------------------------------------------------------
           test4
         5 nexthop                            2000.abcd::cccc.dddd
         6 nexthop                            1000.abcd::1234.5678
         7 interface                          null

switch(config)# pbr-action-list test4 resequence 1 10

           Name
  Sequence Type                              Address/Interface
-------------------------------------------------------------------
           test4
         1 nexthop                            2000.abcd::cccc.dddd
        11 nexthop                            1000.abcd::1234.5678
        21 interface                          null
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

pbr-action-list reset

`pbr-action-list <ACTION-LIST-NAME> reset`

Description

Resets a specified PBR action list to its last successful configuration.

Parameter	Description
<ACTION-LIST-NAME>	Specifies the action list name to be reset.

Examples

```
switch(config)# pbr-action-list test reset
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400	`config`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300 9300S 10000		

policy

```
policy <POLICY-NAME>

    [<SEQUENCE-NUMBER>]
      class {ip|ipv6|mac} <CLASS-NAME>
      action {<REMARK-ACTIONS> | <POLICE-ACTIONS> | <OTHER-ACTIONS>}
      [{<REMARK-ACTIONS> | <POLICE-ACTIONS> | <OTHER-ACTIONS>}]

    [<SEQUENCE-NUMBER>]
    comment ...

no policy <POLICY-NAME>
```

Description

Creates, modifies, or deletes a classifier policy. A policy contains one or more policy entries ordered and prioritized by sequence numbers. Each entry has an IPv4/IPv6/MAC class and one or more policy actions associated with it. An applied policy processes a packet sequentially against policy entries in the list until the last entry in the list has been evaluated or the packet matches an entry. If a match occurs the related entry, actions are taken.

The **no** form of this command is used to delete a policy or an individual policy entry.

Parameter	Description
<POLICY-NAME>	Specifies the name of the policy.
<SEQUENCE-NUMBER>	Specifies a sequence number for the policy entry. Optional. Range: 1 to 4294967295.
`comment`	Stores the remaining entered text as a policy entry comment.
`class {ip|ipv6|mac} <CLASS-NAME>`	Specifies a type of class, **ip** for IPv4, **ipv6** for IPv6 and **mac** for a MAC policy. And specifies a class name.
<REMARK-ACTIONS>	Remark actions can be any of the following options: **{pbr <ACTION-LIST> | pcp <PRIORITY> | ip-precedence <IP-PRECEDENCE-VALUE> | dscp <DSCP-VALUE> | local-priority <LOCAL-PRIORITY-VALUE>}** where: pbr <*ACTION-LIST*> Specifies the PBR action list to be used. pcp <*PCP-VALUE*> Specifies Priority Code Point (PCP) value. Range: 0 to 7.

Parameter	Description
	ip-precedence *<IP-PRECEDENCE-VALUE>* Specifies the numeric IP precedence value. Range: 0 to 7. dscp *<DSCP-VALUE>* Specifies a Differentiated Services Code Point (DSCP) value. Enter either a numeric value (0 to 63) or a keyword as follows: **AF11** - DSCP 10 (Assured Forwarding Class 1, low drop probability) **AF12** - DSCP 12 (Assured Forwarding Class 1, medium drop probability) **AF13** - DSCP 14 (Assured Forwarding Class 1, high drop probability) **AF21** - DSCP 18 (Assured Forwarding Class 2, low drop probability) **AF22** - DSCP 20 (Assured Forwarding Class 2, medium drop probability) **AF23** - DSCP 22 (Assured Forwarding Class 2, high drop probability) **AF31** - DSCP 26 (Assured Forwarding Class 3, low drop probability) **AF32** - DSCP 28 (Assured Forwarding Class 3, medium drop probability) **AF33** - DSCP 30 (Assured Forwarding Class 3, high drop probability) **AF41** - DSCP 34 (Assured Forwarding Class 4, low drop probability) **AF42** - DSCP 36 (Assured Forwarding Class 4, medium drop probability) **AF43** - DSCP 38 (Assured Forwarding Class 4, high drop probability) **CS0** - DSCP 0 (Class Selector 0: Default) **CS1** - DSCP 8 (Class Selector 1: Scavenger) **CS2** - DSCP 16 (Class Selector 2: OAM) **CS3** - DSCP 24 (Class Selector 3: Signaling) **CS4** - DSCP 32 (Class Selector 4: Real time) **CS5** - DSCP 40 (Class Selector 5: Broadcast video) **CS6** - DSCP 48 (Class Selector 6: Network control) **CS7** - DSCP 56 (Class Selector 7) **EF** - DSCP 46 (Expedited Forwarding) `local-priority <LOCAL-PRIORITY-VALUE>` Specifies a local priority value. Range: 0 to 7.

<POLICE-ACTIONS>	Police actions can be the following **{cir <RATE-BPS> cbs <BYTES> exceed}** where:
	`cir <RATE-BPS>`
	Specifies a Committed Information Rate value in Kilobits per second. Range: 1 to 4294967295.
	`cbs <BYTES>`
	Specifies a Committed Burst Size value in bytes. Range: 1 to 4294967295.
	`exceed`
	Specifies action to take on packets that exceed the rate limit.

Parameter	Description
<OTHER-ACTIONS>	Other actions can be the following: `drop` Specifies drop traffic.

Restrictions

MAC classes are not applicable to policies containing PBR actions. Applying such policies to an interface are blocked.

Usage

- For Policy Based Routing, the policy action keyword is **pbr** which itself takes the name of a PBR action list as a parameter.
- A policy entry that contains a PBR action can contain other action types as well.
- An applied policy processes a packet sequentially against policy entries in the list until the last policy entry in the list has been evaluated or the packet matches an entry.
- Entering an existing **<POLICY-NAME>** value will cause the existing policy to be modified, with any new **<SEQUENCE-NUMBER>** value creating an additional policy entry, and any existing **<SEQUENCE-NUMBER>** value replacing the existing policy entry with the same sequence number.
- If no sequence number is specified, a new policy entry is appended to the end of the entry list with a sequence number equal to the highest policy entry currently in the list plus 10.

Examples

Create a policy with two PBR actions:

```
switch(config)# policy pbr_policy
switch (config-policy)# 10 class ip v4_class action pbr action_list1
switch (config-policy)# 20 class ipv6 v6_class action pbr action_list2
switch (config-policy)# exit
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360	`config` The `policy` command takes you into the `config-policy` context where you enter the policy entries.	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
9300 9300S 10000		

primary secondary nexthop

```
primary <PRIMARY-NEXTHOP-IP-ADDRESS> <PRIMARY-NEXTHOP-EGRESS-INTERFACE>
secondary <SECONDARY-NEXTHOP-IP-ADDRESS> <SECONDARY-NEXTHOP-EGRESS-INTERFACE>

no primary <PRIMARY-NEXTHOP-IP-ADDRESS> <PRIMARY-NEXTHOP-EGRESS-INTERFACE>
no secondary <SECONDARY-NEXTHOP-IP-ADDRESS> <SECONDARY-NEXTHOP-EGRESS-INTERFACE>
```

Description

Configures primary and secondary nexthops for the backup nexthop group.

The **no** form of this command removes primary or secondary nexthop for a backup nexthop group.

Parameter	Description
<PRIMARY-NEXTHOP-IP-ADDRESS>	Specifies primary nexthop's IP address.
<PRIMARY-NEXTHOP-EGRESS-INTERFACE>	Specifies primary nexthop's egress interface.
<SECONDARY-NEXTHOP-IP-ADDRESS>	Specifies secondary nexthop's IP address
<SECONDARY-NEXTHOP-EGRESS-ADDRESS>	Specifies secondary nexthop's egress interface.

Usage

- Only one primary and one secondary nexthop per backup nexthop group.
- Configuring more than once will overwrite the previous configuration.
- Primary and secondary nexthop must be configured in the backup nexthop group for the group to be applied.
- Primary and secondary nexthop egress interface must have routing enabled.

Example

Configuring primary and secondary nexthops for a backup nexthop group:

```
switch(config-nexthop-group-grp1)# primary 100.0.0.5 1/1/1
switch(config-nexthop-group-grp1)# secondary 101.0.0.5 1/1/2
```

Configuring primary nexthop for a backup nexthop group with no secondary nexthop configured:

```
switch(config-nexthop-group-grp1)# primary 100.0.0.5 1/1/1
switch(config-nexthop-group-grp1)# secondary 101.0.0.5 1/1/2
switch(config-nexthop-group-grp1)# no secondary 101.0.0.5 1/1/2
```

Command History

Release	Modification
10.11	Command introduced

Command Information

Platforms	Command context	Authority
9300 9300S	`config`	Administrators or local user group members with execution rights for this command.

show ip nexthop group name type backup

`show ip nexthop group <GROUP-NAME> type backup`

Description

Shows configuration for backup nexthop group with a given name.

Parameter	Description
<GROUP-NAME>	Specifies nexthop group name. Nexthop group name can be a maximum of 64 alphanumeric characters.

Example

Showing backup nexthop group configuration:

```
switch# show ip nexthop group grp1 type backup
--------------
Nexthop Group: grp1
   Type:                backup
   Primary Nexthop:     100.0.0.5 1/1/1
   Secondary Nexthop:   101.0.0.5 1/1/2
   Config Status:       Valid
   Active Nexthop:      <Primary|Secondary|None>
```

Command History

Release	Modification
10.11	Command introduced

Command Information

Platforms	Command context	Authority
9300 9300S	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

show ip nexthop group type backup

`show ip nexthop group type backup`

Description

Shows configuration for all backup nexthop groups.

Example

Showing configuration for all backup nexthop groups:

```
switch# show ip nexthop group type backup
--------------
Nexthop Group: grp1
   Type:                backup
   Primary Nexthop:     100.0.0.5 1/1/1
   Secondary Nexthop:   101.0.0.5 1/1/2
   Config Status:       Valid
   Active Nexthop:      <Primary|Secondary|None>
--------------
Nexthop Group: grp2
   Type:                backup
   Primary Nexthop:     200.0.0.5 1/1/3
   Secondary Nexthop:   201.0.0.5 1/1/4
   Config Status:       Valid
   Active Nexthop:      <Primary|Secondary|None>
--------------
Nexthop Group: grp3
   Type:                backup
   Primary Nexthop:     --
   Secondary Nexthop:   101.0.0.5 1/1/2
   Config Status:       Incomplete
--------------
Nexthop Group: grp4
   Type:                backup
   Primary Nexthop:     100.0.0.5 1/1/1
   Secondary Nexthop:   --
   Config Status:       Incomplete
```

Command History

Release	Modification
10.11	Command introduced

Command Information

Platforms	Command context	Authority
9300 9300S	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

show pbr

`show pbr {interface <INTERFACE-NAME>|vrf <VRF-NAME>|summary}`

Description

Shows a detailed view of Policy Based Routing (PBR) in the system.

Parameter	Description
<VRF-NAME>	Specifies name of a VRF.
<INTERFACE-NAME>	Specifies an interface. Format: **member/slot/port**.

Usage

Show commands can only reference the default VRF.

Examples

Showing PBR summary information when there is no active next-hop in the system:

```
switch# show pbr summary
VRF      Port   Policy    PBR      Seq  Type       Nexthop
-----------------------------------------------------------------------
No active PBR nexthop found
-----------------------------------------------------------------------
```

Showing PBR summary information when there are active next-hops in the system:

```
switch# show pbr summary

VRF      Port   Policy    PBR      Seq  Type       Nexthop
-----------------------------------------------------------------------
default  1/1/1  policy_1  pbr_1    10   nexthop    1.1.1.1 (active)
         1/1/2  policy_2  pbr_2    20   nexthop    5.5.5.5 (active)
-----------------------------------------------------------------------
```

Showing PBR summary information when displaying a policy with a **pbr-action-list** applied on a VxLAN L3VNI:

```
switch# configure terminal
switch(config)# interface vxlan 1
switch(config-vxlan-if)# vni 10000
switch(config-vni-10000)# apply policy p1 routed-in
switch(config-vni-10000)# show pbr summary
VRF
      Port
                Policy
                          Class
                                    PBR
                                           Sequence  Type          Nexthop
-------------------------------------------------------------------------------
red
      vni10000
                p1
                          c1
                                    pbr1
                                                10  nexthop       11.2.1.4
(active)
-------------------------------------------------------------------------------
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show pbr-action-list

```
show pbr-action-list [<ACTION-LIST-NAME>] [commands] [configuration] [vsx-peer]
```

Description

Shows the current PBR action list configuration. Action list entries are displayed in ascending order of their sequence number.

Parameter	Description
<ACTION-LIST-NAME>	Specifies the PBR action list name.
`commands`	Formats output as CLI commands.
`configuration`	Displays user-specified configuration.
`vsx-peer`	Displays VSX peer switch information.

Restrictions

If an action list entry is modified to an invalid value (for example through the REST interface), this command will indicate a mismatch for that action entry when run. In this event, use the **pbr-action-list <NAME> reset** command to restore it to the previous valid value.

Usage

- This command does not indicate whether the action list is configured in a policy or applied to an interface. Use the **show pbr** command for PBR status involving action lists.
- A single action list is shown by specifying its name or you can show all action lists by omitting a name argument.
- Using the additional commands keyword, you can change the tabulated output to a configuration style output for single or all list display.

Examples

Create two PBR action lists then run **show pbr-action-list** to display all configured action lists in the default configuration mode:

```
switch(config)# pbr-action-list v4_pbr
switch(config-pbr-action-list-v4_pbr)# 1 nexthop 1.1.1.1
switch(config-pbr-action-list-v4_pbr)# 5 default-nexthop 2.2.2.2
switch(config-pbr-action-list-v4_pbr)# 10 interface null
switch(config-pbr-action-list-v4_pbr)# exit
switch(config)#
switch(config)# pbr-action-list v6_pbr
switch(config-pbr-action-list-v6_pbr)# 20 nexthop 2000:abcd::cccc:dddd
switch(config-pbr-action-list-v6_pbr)# 40 default-nexthop 1000:abcd::1234:5678
switch(config-pbr-action-list-v6_pbr)# 60 interface null
switch(config-pbr-action-list-v6_pbr)# exit
switch#

switch# show pbr-action-list
           Name
           Additional PBR-Action-List Parameters
Sequence     Type               Nexthop
-------------------------------------------------------------------------------
-----
             v4_pbr
1            nexthop            1.1.1.1
5            default-nexthop    2.2.2.2
10           interface          null

             v6_pbr
20           nexthop            2000:abcd::cccc:dddd
40           default-nexthop    1000:abcd::1234:5678
60           interface          null
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	

show running-config current-context

`show running-config current-context`

Description

Displays the configuration of the PBR action list in the current configuration context, in commands mode.

Parameter	Description
`running-config`	Shows configuration currently running on switch.
`current-context`	Limits display to current config context only, in commands mode.

Usage

Useful for reexamining entries previously entered into the action list after its entries have scrolled off the terminal due to other output or upon reentering the context of an existing action list.

Examples

Creating two PBR action lists and running **show running-configuration curent-context** to display the action list configuration in commands mode:

```
switch(config)# pbr-action-list v4_pbr
switch(config-pbr-action-list-v4_pbr)# 1 nexthop 1.1.1.1
switch(config-pbr-action-list-v4_pbr)# 5 default-nexthop 2.2.2.2
switch(config-pbr-action-list-v4_pbr)# 10 interface null
switch(config-pbr-action-list-v4_pbr)# exit
switch(config)#
switch(config)# pbr-action-list v6_pbr
switch(config-pbr-action-list-v6_pbr)# 20 nexthop 2000:abcd::cccc:dddd
switch(config-pbr-action-list-v6_pbr)# 40 default-nexthop 1000:abcd::1234:5678
switch(config-pbr-action-list-v6_pbr)# 60 interface null
switch(config-pbr-action-list-v6_pbr)#
switch(config-pbr-action-list-v6_pbr)# show running-config current-context
pbr-action-list v6_pbr
    20 nexthop 2000:abcd::cccc:dddd
    40 default-nexthop 1000:abcd::1234:5678
    60 interface null
```

Switching context back to the first actionl ist and running the same command:

```
switch(config-pbr-action-list-v6_pbr)# pbr-action-list v4_pbr
switch(config-pbr-action-list-v4_pbr)#
switch(config-pbr-action-list-v4_pbr)# show running-config current-context
pbr-action-list v4_pbr
    1 nexthop 1.1.1.1
    5 default-nexthop 2.2.2.2
    10 interface null
```

Removing action list entry number 5 and running the command again:

```
switch(config-pbr-action-list-v4_pbr)# no 5
switch(config-pbr-action-list-v4_pbr)# show running-config current-context
pbr-action-list v4_pbr
    1 nexthop 1.1.1.1
    10 interface null
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

Chapter 13
IP Directed Broadcast

IP Directed Broadcast is a feature by which remote administration tasks such as backups and wake-on-LAN (WOL) application can be achieved by sending directed broadcast packets for hosts/servers residing on a different subnet. The datagram is routed by normal mechanisms until it reaches a gateway attached to the destination hardware network, at which point it is broadcast. This class of broadcasting is also known as 'directed broadcasting'. IP Directed Broadcast packets are Layer 3 subnet broadcasts.

For example, for the broadcast IP address 192.168.1.255, the corresponding egress interface Subnet IP would be 192.168.1.x/24. Classless subnetting (CIDR) is also supported. For example, for the broadcast IP address 10.10.15.255, the broadcast address for the corresponding interface IP would be 10.10.15.x/21.

IP Directed Broadcast is supported on Route Only Port (ROP), Switched Virtual Interface (SVI), Layer 3 Link Aggregation Group (L3LAG) and Virtual Extensible Lan (VxLAN) interfaces, but is disabled by default. The feature is supported for the associated interface addresses for both the primary and secondary addresses. IP directed broadcast can be enabled on all the configured L3 interfaces. It is only supported for IPv4 addresses.

- VxLAN is only supported on the 6300, 6400, 8100 and 8360 Switch Series.
- VxLAN interfaces require IP Directed Broadcast to be enabled on a Switched Virtual Interface associated with the VxLAN interface.
- IP Directed Broadcast is not supported with Dynamic VRF Route Leak.

For more information on this feature, see the related video on the HPE Aruba Networking AirHeads Broadcasting Channel.

IP Directed Broadcast configuration example

The following are sample topology diagrams for an IP Directed Broadcast configuration. Figure 1 shows when the egress interface is an SVI, while figure 2 shows the diagram when the egress interface is an ROP.

Figure 1 *SVI: IP Directed Broadcast*

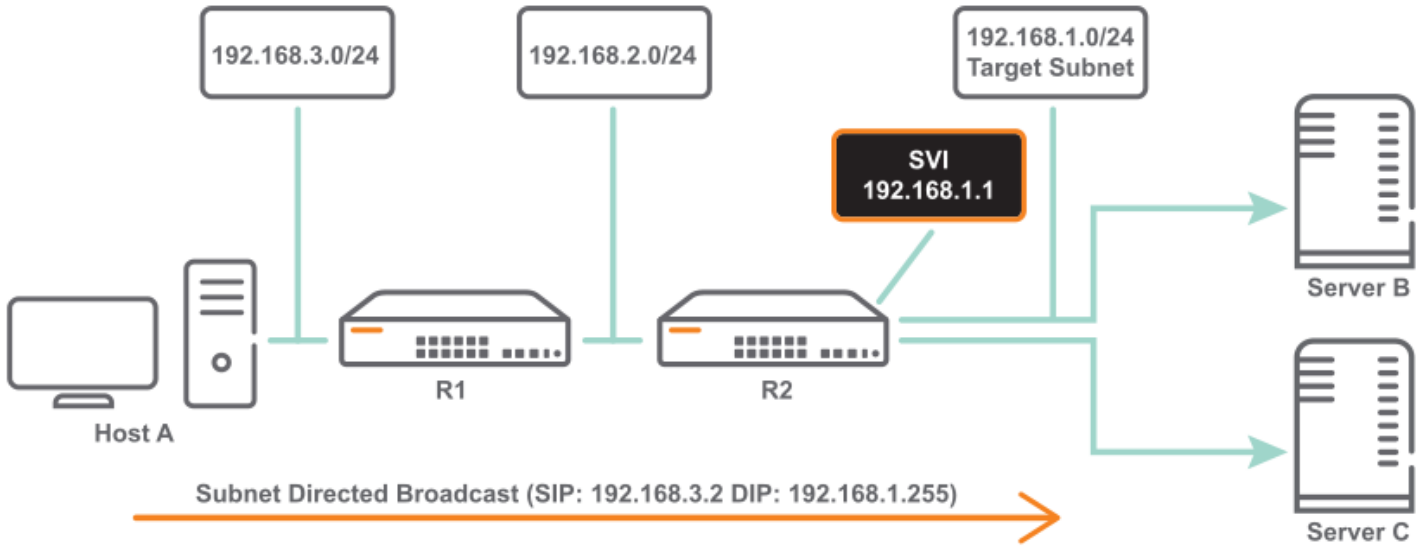

Figure 2 *ROP: IP Directed Broadcast*

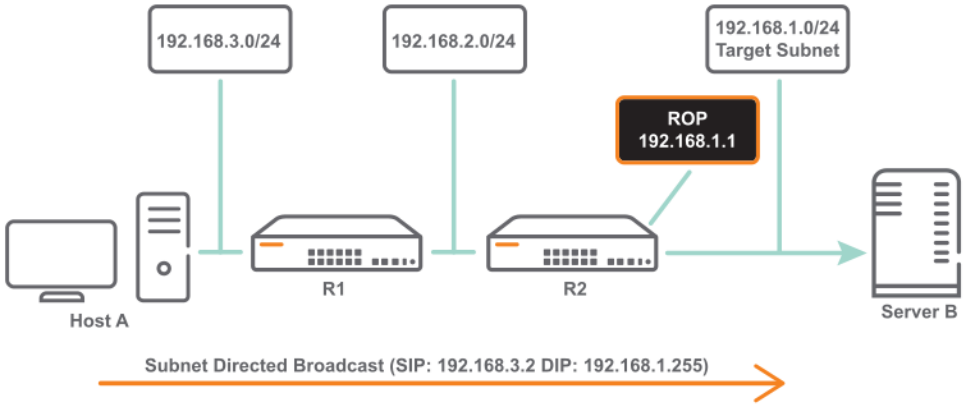

Intermediate routers forward IP Directed Broadcast packets as Unicast. The IP directed broadcast packet is broadcast or flood in the target subnet (DA MAC: All 0xFFs) only after the last hop router.

Host A in subnet 192.168.3.0/24 wants to inject a IP Directed Broadcast (192.168.1.255) packet into Target Subnet 192.168.1.0/24. Router R1 forwards the IP Datagram with DIP 192.168.1.255 as a regular Unicast Datagram. Router R2 then floods the IP Datagram over egress ROP, SVI or VxLAN interface with Destination MAC as all 0xFFs.

At Ingress, Port Based ACLs (PACL) and VLAN Based ACLs (VACL) can be used to restrict/allow IP Directed Broadcast traffic. Existing Port based ACLs (PACL) can be used to allow or disallow certain IP Directed Broadcast Traffic.

An ACL can be configured using the `access-list ip <ACL-NAME>` command and then applied using the `apply access-list ip <ACL-NAME>` command as shown in the following output.

```
switch(config)# access-list ip ipdbacl

switch(config)# interface 1/1/1
switch(config-if)# apply access-list ipdbacl
  in    Inbound (ingress) traffic
```

```
  out  Outbound (egress) traffic

switch(config-if)# int lag 10
switch(config-lag-if)# apply access-list ipdbacl
  in   Inbound (ingress) traffic
  out  Outbound (egress) traffic
```

The following is an example of the `show running-config` command on an ROP interface.

```
switch(config)# interface 1/1/1
    no shutdown
    ip address 192.168.1.1/24
    ip directed-broadcast
```

The following is an example of the `show running-config` command on an SVI interface.

```
switch(config)# vlan 10
interface vlan10
    no shutdown
    ip address 192.168.1.1/24
    ip directed-broadcast
```

The following is an example of the `show running-config` command on an L3LAG interface.

```
switch(config)# interface lag 3
    no shutdown
    ip address 192.168.1.1/24
    ip directed-broadcast
```

Note: Currently egress ACL is supported only on ROP and LAG interfaces, and not on an SVI interface.

IP Directed Broadcast commands

copy support-file feature

`copy support-file feature l3`

Description

Captures support logs to debug any IP Directed Broadcast issues.

IP Directed Broadcast is not supported on subinterfaces on 8100, 6300, 6400, 8325, 8360 and 10000 Switch series.

Examples

Capturing the support logs into a local file:

```
switch# copy support-file feature l3 sftp
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Manager (#)	Administrators or local user group members with execution rights for this command.

ip directed-broadcast

```
ip directed-broadcast
no ip directed-broadcast
```

Description

Turns on IP Directed Broadcast for the specified interface. The **no** form of this command turns it off. This command is disabled by default.

IP Directed Broadcast is not supported on subinterfaces on 6300, 6400, 8100, 8325, 8360 and 10000 Switch series.

Examples

Enabling and disabling IP Directed Broadcast on an physical interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ip directed-broadcast
switch(config-if)# no ip directed-broadcast
```

Enabling and disabling IP Directed Broadcast on a VLAN interface:

```
switch(config)# interface vlan 100
switch(config-if-vlan)# ip directed-broadcast
switch(config-if-vlan)# no ip directed-broadcast
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if`	

show arp

```
show arp
```

Description

Shows IP directed broadcast verification.

IP Directed Broadcast is not supported on subinterfaces on 8100, 6300, 6400, 8325, 8360 and 10000 Switch series.

Examples

Showing IP directed broadcast verification:

```
switch# show arp
IPv4 Address     MAC                Port         Physical Port          State
--------------------------------------------------------------------------------
1.1.1.255        FF:FF:FF:FF:FF:FF  1/1/1        1/1/1                  permanent
3.1.1.255        FF:FF:FF:FF:FF:FF  vlan10                              permanent


Total Number Of ARP Entries Listed: 2.
--------------------------------------------------------------------------------
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config	

show ip interface

```
show ip interface <INTERFACE-NAME>
```

Description

Displays the status of IP Directed Broadcast on the specified interface along with other interface related attributes.

Parameter	Description
<INTERFACE-NAME>	Specifies the interface to use as a source for displaying the status of the IP Directed Broadcast.

Examples

Displaying the IP Directed Broadcast status on the specified interface:

```
switch# show ip interface vlan30


Interface vlan30 is up
 Admin state is up
 Hardware: Ethernet, MAC Address: 94:f1:28:21:63:00
 IP MTU 1500
 IP Directed Broadcast is Enabled
 IPv4 address 192.168.3.1/24
 L3 Counters: Rx Disabled, Tx Disabled

 Statistics                      RX                  TX                Total
 ------------ -------------------- -------------------- --------------------
 L3 Packets                       0                    0                   0
 L3 Bytes                         0                    0                   0
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ip directed-broadcast

```
show ip directed-broadcast
```

Description

Displays the summary of the interfaces on which IP Directed Broadcast is enabled.

Examples

On the HPE Aruba Networking 6400 Switch Series, interface identification differs.

Displaying the summary of the interfaces on which IP Directed Broadcast is enabled:

```
switch# show ip directed-broadcast

 IPv4 Directed Broadcast Configuration

 Interface       Status
 ----------     --------
 1/1/1           Enabled
 vlan10          Enabled
 vlan30          Enabled
```

Displaying IP Directed Broadcast Host entries installed in Neighbor cache:

```
switch# show arp state permanent
IPv4 Address     MAC                Port         Physical Port    State
--------------------------------------------------------------------------------
52.1.1.255       FF:FF:FF:FF:FF:FF  1/1/1        1/1/1            permanent
40.0.0.255       FF:FF:FF:FF:FF:FF  vlan20       vlan20           permanent
Total Number Of ARP Entries Listed- 2.
--------------------------------------------------------------------------------
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

Chapter 14
IP Neighbor Flood

IP Neighbor Flood is only supported on the HPE Aruba Networking 6300, 6400, 8100, 8320, 8325/8325H/8325P, 8360, 9300/9300S and 10000 Switch Series.

When a port goes down, L3 traffic to a neighbor will get dropped until the MAC address is learned on the new port. Source MAC learning could take a while, and this learning could be dependent on traffic being successfully delivered to the client.

An example is when wireless clients are connected through access points balanced across several wireless controllers connected to AOS-CX switches. Traffic to clients is dropped when the port connecting the controller and switch goes down.

The port can go down when the wireless controller goes down and reboots. Access points are adopted by other controllers on the network to avoid traffic drop for the wireless clients, but the back-end data plane requires handling. During such failover conditions in scaled setups, MAC learning could take a while, and significant traffic loss can occur. The solution is to flood traffic on the VLAN so that traffic reaches the hosts. Once neighbors are relearned over a new port on the same VLAN, the switch will stop flooding.

IP Neighbor Flood is supported on SVIs. It is disabled by default.

IP Neighbor Flood commands

ip neighbor-flood

```
ip neighbor-flood
```

Description

Enables VLAN flooding for the specified VLAN interface when a neighbor link goes down.

The **no** form of this command disables VLAN flooding for the specified VLAN interface.

Examples

Enabling IP Neighbor Flood on a VLAN interface:

```
switch(config)# interface vlan 3
switch(config-if-vlan)# ip neighbor-flood
```

Disabling IP Neighbor Flood on a VLAN interface.

```
switch(config)# interface vlan 3
switch(config-if-vlan)# no ip neighbor-flood
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if-vlan`	Administrators or local user group members with execution rights for this command.

show ip interface

```
show ip interface <IFNAME>
```

Description

Displays the status of IP Neighbor Flood on the specified interface along with other interface-related attributes.

Parameter	Description
<IFNAME>	Specifies the interface name (for example, `vlan30`). Optional.

Examples

```
switch# show ip interface vlan30

Interface vlan30 is up
 Admin state is up
 Hardware: Ethernet, MAC Address: 94:f1:28:21:63:00
 IP MTU 1500
 IP Neighbor Flood is Enabled
 IPv4 address 192.168.3.1/24
 L3 Counters: Rx Disabled, Tx Disabled
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ip neighbor-flood

```
show ip neighbor-flood
```

Description

Displays the interfaces on which IP Neighbor Flood is enabled.

Examples

```
switch# show ip neighbor-flood

 IP Neighbor Flood Configuration

 Interface       Status
 ---------       -------
 vlan10          Enabled
 vlan30          Enabled
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show running-config

```
show running-config
```

Description

Displays the current running configuration.

Examples

```
switch# show running-config
interface vlan10
    ip neighbor-flood
interface vlan30
    ip neighbor-flood
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

Chapter 15
Key chain

This chapter provides details for configuring and verifying the functions of a key chain. Key chain configurations are spread across `config` context, `keychain` context, and `keychain-key` context.

A key chain provides seamless authentication-key rollover. When authentication method is enabled, using key chain allows OSPF to overcome static key configuration to change the key periodically. Using key chain also reduces risks of keys being guessed by configuring rotating keys.

A key chain must be used by an application like OSPF that communicates by using the keys with its peers.

Key chain commands

accept-lifetime

```
accept-lifetime [start-time <time> <month>/<day>/<year>] {duration {<seconds> | infinite}
| end-time <time> <month>/<day>/<year>}
```

Description

Configures the duration for which the key is valid for receiving packets.

The **no** form of this command configures the key packet receiving duration to the default value of an infinite time.

Parameter	Description
`start-time`	Time at which the key chain lifetime starts. Required. Format: HH:MM:SS
`end-time`	Time at which the key chain lifetime expires. Required. Format: HH:MM:SS
`day`	Day of the month. Required. Range: 1-31.
`month`	Month of the year. Required.
`year`	Year. Required. Range: 2020-2050
`duration`	Time in seconds. Optional. Range: 1-2147483646.
`infinite`	Specifies infinite time for the key. Optional.

Examples

Configuring the duration for which the key is valid for receiving packets:

```
switch# configure terminal
switch(config)# keychain ospf_keys
switch(config-keychain)# key 1
switch(config-keychain-key)# accept-lifetime start-time 10:10:10 10/25/2020 end-
time 10:10:10 11/25/2020
switch(config-keychain-key)# accept-lifetime start-time 10:10:10 10/25/2020
duration 1000
switch(config-keychain-key)# accept-lifetime start-time 10:10:10 10/25/2020
duration infinite
switch(config-keychain-key)# accept-lifetime end-time 10:10:10 11/25/2020
switch(config-keychain-key)# accept-lifetime duration 1000
switch(config-keychain-key)# accept-lifetime duration infinite
```

Configuring the key packet receiving duration to the default value of an infinite time:

```
switch# configure terminal
switch(config)# keychain ospf_keys
switch(config-keychain)# key 1
switch(config-keychain-key)# no accept-lifetime
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	`config-keychain-key`	

cryptographic-algorithm

```
recv-id cryptographic-algorithm {aes-cmac-128 |hmac-sha-1 | hmac-sha-256 | hmac-sha-384 |
hmac-sha-512 | md5}
no cryptographic-algorithm
```

Description

Configures the recv-id cryptographic algorithm for the key. The key will not be valid until the receive ID, the send ID, and send lifetime is configured for TCP-AO, Choose one of the authentication algorithms from the following parameters. The **no** form of this command configures the default cryptographic algorithm for a key, md5.

TCP Authentication Option (TCP-AO) authentication supports only the *aes-cmac-128* and *hmac-sha-1* algorithms. If you are configuring TCP-AO, you must select one of these options.

Parameter	Description
`aes-cmac-128`	Sets the authentication algorithm for the key to AES-CMAC-128. This parameter is only supported for TCP-AO.

Parameter	Description
`hmac-sha-1`	Sets the authentication algorithm for the key to SHA-1. This parameter is also supported for TCP-AO.
`hmac-sha-256`	Sets the authentication algorithm for the key to SHA-256.
`hmac-sha-384`	Sets the authentication algorithm for the key to SHA-384.
`hmac-sha-512`	Sets the authentication algorithm for the key to SHA-512.
`md5`	Sets the authentication algorithm for the key to md5. Maximum length of the key string supported: 16 bytes (**config**-if context), 64 bytes (**config-keychain-key context**).

Examples

Set the authentication algorithm for the key to **SHA-384**:

```
switch(config)# keychain ospf_keys
switch(config-keychain)# key 1
switch(config-keychain-key)# recv-id cryptographic-algorithm hmac-sha-384
```

Set the authentication algorithm to the default, **md5**:

```
switch(config)# keychain ospf_keys
switch(config-keychain)# key 1
switch(config-keychain-key)# no recv-id cryptographic-algorithm
```

Command History

Release	Modification
10.11	The **aes-cmac-128** parameter is introduced.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-keychain-key`	Administrators or local user group members with execution rights for this command.

key

`key <KEY-ID>`

Description

Creates the key for a key chain and enters the key chain key context. A maximum of 64 keys can be configured per key chain.

The **no** form of this command deletes the key from the key chain.

Parameter	Description
<KEY-ID>	ID of the key. Required. Range: 1-255.

Examples

Creating a key for a key chain:

```
switch# configure terminal
switch(config)# keychain ospf_keys
switch(config-keychain)# key 1
```

Deleting a key from a key chain:

```
switch# configure terminal
switch(config)# keychain ospf_keys
switch(config-keychain)# no key 1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	`config-keychain`	

keychain

`keychain <KEYCHAIN-NAME>`

Description

Creates the key chain and enters the key chain context. A maximum of 64 key chains can be configured in the system.

The **no** form of this command removes the key chain if it is not used by any subscribers.

Parameter	Description
<KEYCHAIN-NAME>	Name of the key chain. Required.

Examples

Creating a key chain:

```
switch# configure terminal
switch(config)# keychain ospf_keys
```

Removing a key chain:

```
switch# configure terminal
switch(config)# keychain ospf_keys
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	config	

key-string

`key-string [{ciphertext | plaintext} <PASSWORD>]`

Description

Sets the key password. The password is internally stored in encrypted form. The key is not valid until its password has been set.

The **no** form of this command deletes the password used for the key.

Parameter	Description
ciphertext	Specifies that the key password is provided as ciphertext.
plaintext	Specifies that the key password is provided as plaintext.
<PASSWORD>	Specifies the key password.

When the key password is not provided on the command line, plaintext password prompting occurs upon pressing Enter. The entered password characters are masked with asterisks.

Examples

Setting the key password with plaintext:

```
switch(config)# keychain ospf_keys
switch(config-keychain)# key 1
switch(config-keychain-key)# key-string plaintext F82#450bHP
```

Setting the key password with plaintext prompting:

```
switch(config)# keychain ospf_keys
switch(config-keychain)# key 1
switch(config-keychain-key)# key-string
Enter the key password: *************
Re-Enter the key password: *************
```

Setting the key password with ciphertext:

```
switch(config)# keychain ospf_keys
switch(config-keychain)# key 1
switch(config-keychain-key)# key-string ciphertext AQBpfciFZ/P...biAAAOjc0a8=
```

Deleting the password for the key:

```
switch(config)# keychain ospf_keys
switch(config-keychain)# key 1
switch(config-keychain-key)# no key-string
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	`config-keychain-key`	

name

```
name <KEY-NAME>
no name <KEY-NAME>
```

Description

Configures a name for a numbered key in a key chain.

The **no** form of this command removes the name of the key.

Parameter	Description
<KEY-NAME>	Specifies the name of the key in alphanumeric characters. Range: 1-64.

Examples

Creating a name for a key in a key chain called **abcdef123456**:

```
switch# configure terminal
switch(config)# keychain macsec_keys
```

```
switch(config-keychain)# key 1
switch(config-keychain-key)# name abcdef123456
```

Removing the name of the key named **abcdef123456**:

```
switch# configure terminal
switch(config)# keychain macsec_keys
switch(config-keychain)# key 1
switch(config-keychain-key)# no name abcdef123456
```

Command History

Release	Modification
10.11	Command added.

Command Information

Platforms	Command context	Authority
All platforms	`config-keychain-key`	

recv-id

`recv-id <0-255>`

Description

Configures the receive ID for a keychain key. The receive ID has to be unique across keys in the keychain.

The **no** form of this command configures removes the recv-id value. The receive ID can not be changed for an active key of a keychain which is associated with BGP neighbor.

Parameter	Description
`<0-255>`	Set the receive ID corresponding to the keychain key. Supported values are 0-255.

Examples

Configuring the receive ID for the keychain key.

```
switch# configure terminal
switch(config)# keychain ospf_keys
switch(config-keychain)# key 1
switch(config-keychain-key)# recv-id 1
```

Command History

Release	Modification
10.11	Command introduced.

Command Information

Platforms	Command context	Authority
All platforms	`config-keychain-key`	

send-id

```
send-id <0-255>
```

Description

Configures the send ID for a keychain key. The send ID has to be unique across keys in the keychain.

The **no** form of this command configures removes the send-id value. The send id can not be changed for an active key of a keychain which is associated with BGP neighbor.

Parameter	Description
`<0-255>`	Set the send IDcorresponding to the keychain key. Supported values are 0-255.

Examples

Configuring the send ID for the keychain key.

```
switch# configure terminal
switch(config)# keychain ospf_keys
switch(config-keychain)# key 1
switch(config-keychain-key)# send-id 218
```

Command History

Release	Modification
10.11	Command introduced.

Command Information

Platforms	Command context	Authority
All platforms	`config-keychain-key`	

send-lifetime

```
send-lifetime [start-time <time> <month>/<day>/<year>] {duration {<seconds> | infinite} |
end-time <time> <month>/<day>/<year>}
```

Description

Configures the duration for which the key is valid for sending packets.

The **no** form of this command configures the key packet sending duration to the default value of an infinite time.

Parameter	Description
`start-time`	Time at which the key chain lifetime starts. Required. Format: HH:MM:SS
`end-time`	Time at which the key chain lifetime expires. Required. Format: HH:MM:SS
`day`	Day of the month. Required. Range: 1-31.
`month`	Month of the year. Required.
`year`	Year. Required. Range: 2020-2050
`duration`	Time in seconds. Optional. Range: 1-2147483646.
`infinite`	Specifies infinite time for the key. Optional.

Examples

Configuring the duration for which the key is valid for sending packets:

```
switch# configure terminal
switch(config)# keychain ospf_keys
switch(config-keychain)# key 1
switch(config-keychain-key)# send-lifetime start-time 10:10:10 10/25/2020 end-time
10:10:10 11/25/2020
switch(config-keychain-key)# send-lifetime start-time 10:10:10 10/25/2020 duration
1000
switch(config-keychain-key)# send-lifetime start-time 10:10:10 10/25/2020 duration
infinite
switch(config-keychain-key)# send-lifetime end-time 10:10:10 11/25/2020
switch(config-keychain-key)# send-lifetime duration 1000
switch(config-keychain-key)# send-lifetime duration infinite
```

Configuring the key packet sending duration to the default value of an infinite time:

```
switch# configure terminal
switch(config)# keychain ospf_keys
switch(config-keychain)# key 1
switch(config-keychain-key)# no send-lifetime
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	`config-keychain-key`	

show capacities keychain

```
show capacities keychain
```

Description

Shows the maximum number of key chains and keys configurable in a key chain.

Example

```
switch# show capacities keychain

System Capacities: Filter Keychain
Capacities Name
                   Value
-------------------------------------------------------------------------------
-------------------------
Maximum number of keychains supported in the system
                    64
Maximum number of Keys supported in a single Keychain
                    64
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Operator (>)	Administrators or local user group members with execution rights for this command.

show keychain

```
show keychain [<KEYCHAIN-NAME>]
```

Description

Shows information about configured and active keys of a named key chain or (if **keychain-name** is not specified) all configured key chains.

Parameter	Description
<KEYCHAIN-NAME>	Name of the key chain. Optional.

Example

```
switch# show keychain
Keychain Name : macsec_keys
  Number of Keys      : 1
  Active Send Key ID  :
  Active Recv Key IDs :

  Key ID : 1
```

```
    Key name          : abcdef123456
    Key string        :
AQBapYa+0qQDzcakbB1TopeX0AMYDDWDW015orkH5mY3qJDaBAAAADASiBQ=
    Send Key Validity : 00:00:00 01/01/2020 to Infinite
    Recv Key Validity : 00:00:00 01/01/2020 to Infinite

Keychain Name : ospf_keys
  Number of Keys      : 2
  Active Send Key ID  : 7
  Active Recv Key IDs : 7, 200

  Key ID              : 7
    Key name          : -
    Key string        :
AQBapZ1OHiO9W3JwRqnjtLfbV73BPLS1S6TGVg+Lzl7N4e5eBAAAAPWaPBE=
    Crypto-Algorithm  : sha256
    Send Key Validity : 00:00:01 10/1/2020 to 23:59:01 10/1/2021
    Recv Key Validity : 00:00:01 10/1/2020 to infinite
  Key ID              : 200
    Key name          : -
    Key string        :
AQBapZ1OHiO9W3JwRqnjtLfbV73BPLS1S6TGVg+Lzl7N4e5eBAAAAPWaPBE=
    Crypto-Algorithm  : sha512
    Send Key Validity : 00:00:01 10/1/2020 to 23:59:01 10/1/2021
    Recv Key Validity : 00:00:01 10/1/2020 to 23:59:01 10/1/2021

Keychain Name : bgp_keys
  Number of Keys      : 2
  Active Send Key ID  : 7
  Active Recv Key IDs : 7

  Key ID              : 7
    Key name          : -
    Key string        :
AQBapZ1OHiO9W3JwRqnjtLfbV73BPLS1S6TGVg+Lzl7N4e5eBAAAAPWaPBE=
    Crypto-Algorithm  : md5
    Send Key Validity : 00:00:01 10/26/2020 to 23:59:01 10/1/2021
    Recv Key Validity : 00:00:01 10/22/2020 to infinite
  Key ID              : 8
    Key name          : -
    Key string        :
AQBapZ1OHiO9W3JwRqnjtLfbV73BPLS1S6TGVg+Lzl7N4e5eBAAAAPWaPBE=
    Crypto-Algorithm  : sha384
    Send Key Validity : 00:00:01 10/1/2021 to 23:59:01 10/1/2021
    Recv Key Validity : 00:00:01 10/1/2021 to 23:59:01 10/1/2021
```
```
Keychain Name : ospf_keys
  Number of Keys      : 2
  Active Send Key ID  : 7
  Active Recv Key IDs : 7, 200

  Key ID              : 7
    Key name          : -
    Key string        :
AQBapZ1OHiO9W3JwRqnjtLfbV73BPLS1S6TGVg+Lzl7N4e5eBAAAAPWaPBE=
    Crypto-Algorithm  : sha256
    Send Key Validity : 00:00:01 10/1/2020 to 23:59:01 10/1/2021
    Recv Key Validity : 00:00:01 10/1/2020 to infinite
  Key ID              : 200
    Key name          : -
    Key string        :
```

```
AQBapZ1OHiO9W3JwRqnjtLfbV73BPLS1S6TGVg+Lzl7N4e5eBAAAAPWaPBE=
    Crypto-Algorithm  : sha512
    Send Key Validity : 00:00:01 10/1/2020 to 23:59:01 10/1/2021
    Recv Key Validity : 00:00:01 10/1/2020 to 23:59:01 10/1/2021
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Operator (>)	Administrators or local user group members with execution rights for this command.

show running-config keychain

```
show runnning-config keychain
```

Description

Shows the configurations for key chain protocol.

Example

```
switch# show running-config keychain
keychain ospf_keys
  key 1
    key-string ciphertext
AQBapZ1OHiO9W3JwRqnjtLfbV73BPLS1S6TGVg+Lzl7N4e5eBAAAAPWaPBE=
    cryptographic-algorithm md5
    accept-lifetime start-time 10:10:10 10/25/2020 end-time 10:10:10 11/25/2020
    send-lifetime start-time 10:10:10 10/25/2020 end-time 10:10:10 11/25/2020
  key 45
    key-string ciphertext
AQBapZ1OHiO9W3JwRqnjtLfbV73BPLS1S6TGVg+Lzl7N4e5eBAAAAPWaPBE=
    accept-lifetime start-time 10:10:10 10/25/2020 end-time 10:10:10 11/25/2020
  key 33
keychain macsec_keys
  key 1
    name abcdef123456
    key-string ciphertext
AQBapYa+0qQDzcakbB1TopeX0AMYDDWDW015orkH5mY3qJDaBAAAADASiBQ=
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Operator (>)	Administrators or local user group members with execution rights for this command.

Chapter 16
IP Client Tracker

IP Client Tracker is only supported on the HPE Aruba Networking 6x00 Switch Series.

The client IP address tracking feature will learn and update the IP addresses of the access devices and clients connected to the switch. It can track addresses of directly connected clients, as well as clients connected to a downstream device such as a wireless access point.

IP Client tracker can be enabled for user-based tunneling (UBT) clients to learn the IP addresses on the required VLANs and interfaces.

Enabling IP Client Tracker based on UBT modes:

- If the UBT mode is Local VLAN, enable IP Client tracker for the UBT client VLAN defined using the `ubt-client-vlan` command.
- If the UBT mode is VLAN extend, enable IP Client tracker for the UBT client VLAN defined under role.

client track ip

`client track ip`

Description

Enables client IP address tracking on the switch. The default is disabled on global and VLAN levels.

Admin users can enable client IP address tracking at the VLAN level.

Tracking enabling will take effect only if the client IP address tracking is enabled at system and VLAN level.

The **no** form of the command disables client IP address tracking. If tracking is disabled at switch level, it will be stopped even if it is enabled at VLAN or port level.

Example

Enable client IP address tracking at switch level:

```
switch(config)# client track ip
```

Enable client IP address tracking on VLAN 100:

```
switch(config)# vlan 100
switch(config-vlan-100)# client track ip
```

Enable client IP address tracking on VLANs 10 to 100:

```
switch(config)# vlan 10-100
switch(config-vlan-<10-100>)# client track ip
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400	config	

client track ip { enable | disable | auto }

```
client track ip { enable | disable | auto }
```

Description

Enables client IP address tracking on the specified set of interfaces. Tracking will take effect only if client IP address tracking is enabled at both the system level and for the VLAN to which the port belongs. Default: auto.

The **no** form of the command disables client IP address tracking on the specified set of interfaces.

Parameter	Description
enable	Specifies that all client IP addresses will be tracked in the port.
disable	Specifies that client IP addresses will not be tracked in the port.
auto	Specifies the following: For LLDP devices: Only the specified client IP address will be tracked in the port and other client IP addresses will not be tracked. For non-LLDP devices: All client IP addresses will be tracked in the port.

Example

Enable client IP address tracking on interface 1/1/1:

```
switch(config)# interface 1/1/1
switch(config-if)# client track ip enable
```

Enable client IP address tracking on interfaces 1/1/1 to 1/1/5:

```
switch(config)# interface 1/1/1-1/1/5
switch(config-if-<1/1/1-1/1/5>)# client track ip enable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400	config-if	

client track ip client-limit

```
client track ip client-limit <CLIENT-LIMIT>
```

Description

Configures the maximum number of clients to be tracked on the specified set of interfaces.

The **no** form of the command resets the client limit to the default value. Default values vary according to switch model:

- 6300: 2048
- 6400: 4096

Parameter	Description
CLIENT-LIMIT	Specifies the maximum number of clients tracked on a port. Required. Range: 1-2048 (6300) 1-4096 (6400). Default: 2048 (6300) 4096 (6400)..

Example

Configure the maximum number of clients to be tracked on interface 1/1/5:

```
switch(config)# interface 1/1/5
switch(config-if)# client track ip client-limit 32
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400	config-if	

client track ip update-interval

```
client track ip update-interval <INTERVAL>
```

Description

Configures how often client IP addresses are updated.

The **no** form of the command resets the update interval to the default of 1800 seconds.

Parameter	Description
INTERVAL	Specifies the update interval in seconds. Required. Range: 60-28000. Default: 1800.

Example

Configure the update interval for an interface:

```
switch(config)# interface 1/1/1
switch(config-if)# client track ip update-interval 600
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400	config-if	

client track ip update-method probe

```
client track ip update-method probe
```

Description

Enables probing the client to update the IP address.

The probe is sent to all clients on the tracking list that have an IP address in the following scenarios:

1. IP packets are not received from the clients during the IP address update cycle.
2. There is no IP packet from a learned IP address. In this case, a probe will be sent for the IP address to confirm if it is still owned by that client.

The **no** form of the command disables probing.

Example

Disable probing to update the client IP address:

```
switch(config)# no client track ip update-method probe
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config	

show capacities

```
show capacities
```

Description

Shows the capacities configured on the switch.

Example

```
switch# show capacities

System Capacities:
Capacities Name                                                                    Value
-----------------------------------------------------------------------------------------
Maximum number of Access Control Entries configurable in a system                14336
Maximum number of Access Control Lists configurable in a system                   1024
Maximum number of class entries configurable in a system                          1024
Maximum number of classes configurable in a system                                 512
Maximum number of entries in an Access Control List                               1024
Maximum number of entries in a class                                              1024
Maximum number of entries in a policy                                             1024
Maximum number of classifier policies configurable in a system                    512
Maximum number of policy entries configurable in a system                        1024
Maximum number of clients supported for tracking the IP address in the system    128

switch# show capacities client-track-ip-client-limit

System Capacities: Filter Client Track IP Client Limit
Capacities Name
Value
-----------------------------------------------------------------------------------------
-
Maximum number of clients supported for tracking the IP address in the system
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>)	Administrators or local user group members with execution rights for this command.

show client ip { count | port | vlan }

```
show client ip { count | port | vlan }
```

Description

Shows number of client IP addresses or information about client IP addresses tracked on ports and VLANs.

Parameter	Description
count	Displays number of clients tracked.
port	Displays client IP addresses tracked on the ports.
vlan	Displays client IP addresses tracked on the VLANs.

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400	Operator (>)	Administrators or local user group members with execution rights for this command.

Chapter 17
Routing Information Protocol (RIP)

Routing Information protocol (RIP, RIPv2, RIPng) is a distance-vector routing protocol that uses hop count as a routing metric. It is useful in small networks. RIP can be configured on L3 ports, LAG, VLAN and Loopback interfaces. All configurations work in the interface context and require the associated interface to be routing interface.

- **RIP/RIPv2:** IPv4 implementation of Routing Information protocol defined in RFC 2453.
- **RIPng:** IPv6 implementation of Routing Information protocol defined in RFC 2080.

Overview

RIP characteristics:

- Prevents routing loops by adding a limit to the number of hops allowed in a path from source to destination.
- Allows a maximum of 15 hops which limits the size of the network that RIP can support.
- Uses broadcast UDP data packets to exchange routing information.
- RIPv2 is a classless protocol that supports variable-length subnet mask (VLSM), classless inter-domain routing (CIDR) and route summarization.
- RIPng uses ipsec for message authentication.

RIP limitations:

- RIPv1 is not supported.
- Only poison reverse is supported for loop prevention.
- RIP/RIPng metric configuration support is not available.
- Redistribute RIP/RIPng is not supported in BGP/BGP+.

RIPv2 (IPv4) commands

Configuration commands

router rip

```
router rip <PROCESS-ID> [vrf <VRF-NAME>]
no router rip <PROCESS-ID> [vrf <VRF-NAME>]
```

Description

Creates RIP process if not already created and enters the **router rip <PROCESS-ID>** context for the VRF mentioned. If no VRF is mentioned, a default is used. Only one RIP process is allowed per VRF.

The **no** form of this command deletes the RIP instance for the VRF. If no VRF is mentioned the default is deleted.

Parameter	Description
<PROCESS-ID>	Specifies name of the RIP process ID. Range: 1-63.
vrf <VRF-NAME>	Specifies VRF name.

Examples

Creating RIP process and naming the VRF:

```
switch(config)# router rip 2 vrf red
```

Deleting RIP process:

```
switch(config)# no router rip 2 vrf red
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config	Administrators or local user group members with execution rights for this command.

Interface commands

ip rip

```
ip rip <PROCESS-ID> {all-ip | ip-address}
no ip rip <PROCESS-ID> {all-ip | ip-address}
```

Description

Enables RIP process on an interface.

The **no** form of this command deletes the RIP process from an interface.

Parameter	Description
`ip rip <PROCESS-ID>`	Specifies RIP process ID. Range: 1-63.
`all-ip`	Specifies RIP for all IP addresses configured on the interface.
`ip-address`	Specifies IP address for RIP on the interface.

Usage

- If an IP address is removed from an interface configured with RIP, all RIP configurations will be removed from the interface.
- If **ip rip *1* all-ip** is configured and a new IP address is added to the interface, RIP configurations will not be applicable for the newly added IP address.

Examples

Configuring RIP for all IP addresses configured on the interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ip rip 1 all-ip
```

Deleting RIP for all IP addresses configured on the interface:

```
switch(config)# interface 1/1/1
switch(config-if)# no ip rip 1 all-ip
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if`	Administrators or local user group members with execution rights for this command.

ip rip all-ip enable

```
ip rip all-ip enable
no ip rip all-ip enable
```

Description

Enables RIP process for all RIP enabled IP addresses configured on interface.

The **no** form of this command disables RIP process on the interface.

Usage

- Default settings allow an interface to receive RIP packets.
- If an IP address is removed from an interface configured with RIP, all RIP configurations will be removed from the interface.

Examples

Enabling RIP process for all RIP enabled IP addresses on interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ip rip all-ip enable
```

Disabling RIP process on interface:

```
switch(config)# interface 1/1/1
switch(config-if)# no ip rip all-ip enable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ip rip all-ip disable

```
ip rip all-ip disable
no ip rip all-ip disable
```

Description

Disables RIP process for all RIP enabled IP addresses configured on the interface.

The **no** form of this command enables RIP process for all RIP enabled IP addresses configured on the interface.

Usage

- Default settings allow an interface to receive RIP packets.
- If an IP address is removed from an interface configured with RIP, all RIP configurations will be removed from the interface.

Examples

Disabling RIP process for all RIP enabled IP addresses on interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ip rip all-ip enable
```

Enabling RIP process for all RIP enabled IP addresses on interface:

```
switch(config)# interface 1/1/1
switch(config-if)# no ip rip all-ip enable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if`	Administrators or local user group members with execution rights for this command.

ip rip all-ip send disable

```
ip rip all-ip send disable
no ip rip all-ip send disable
```

Description

Disables interface from sending RIP packets for all RIP enabled IP addresses.

The **no** form of this command enables interface to send RIP packets for all RIP enabled IP addresses.

Usage

- Default settings allow an interface to send RIP packets.
- If an IP address is removed from an interface configured with RIP, all RIP configurations will be removed from the interface.

Examples

Disabling interface from sending RIP packets for all RIP enabled IP addresses on interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ip rip all-ip send disable
```

Enabling interface to send RIP packets for all RIP enabled IP addresses on interface:

```
switch(config)# interface 1/1/1
switch(config-if)# no ip rip all-ip send disable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

ip rip all-ip receive disable

```
ip rip all-ip receive disable
no ip rip all-ip receive disable
```

Description

Disables interface from receiving RIP packets for all enabled IP addresses.

The **no** form of this command enables interface to receive RIP packets for all RIP enabled IP addresses.

Usage

- Default settings allow an interface to receive RIP packets.
- If an IP address is removed from an interface configured with RIP, all RIP configurations will be removed from the interface.

Examples

Disabling interface from receiving RIP packets for all RIP enabled IP addresses on interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ip rip all-ip receive disable
```

Enabling interface to receive RIP packets for all RIP enabled IP addresses on interface:

```
switch(config)# interface 1/1/1
switch(config-if)# no ip rip all-ip receive disable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

Routing commands

enable

```
enable
no enable
```

Description

Enables RIP process if disabled. By default RIP process is enabled.

The **no** form of this command disables the RIP process.

Examples

Enabling RIP process when disabled:

```
switch(config)# router rip 1
switch(config-rip-1)# enable
```

Disabling RIP process when enabled:

```
switch(config)# router rip 1
switch(config-rip-1)# no enable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config	Administrators or local user group members with execution rights for this command.

disable

```
disable
no disable
```

Description

Disables RIP process.

The **no** form of this command enables the RIP process.

Examples

Disabling RIP process:

```
switch(config)# router rip 1
switch(config-rip-1)# disable
```

Enabling RIP process:

```
switch(config)# router rip 1
switch(config-rip-1)# no disable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

distance

```
distance <DISTANCE>
no distance
```

Description

Configures administrative distance for RIP. Administrative distance is used as criteria to select the best route when multiple protocols have the same route.

The **no** form of this command sets the RIP administrative distance to the default. Default: 120.

Parameter	Description
<DISTANCE>	Specifies RIP administrative distance. Range: 1 to 255.

Examples

Configuring administrative distance for RIP:

```
switch(config)# router rip 1
switch(config-rip-1)# distance 100
```

Setting administrative distance for RIP to default values:

```
switch(config)# router rip 1
switch(config-rip-1)# no distance
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

maximum-paths

```
maximum-paths <MAX-VALUE>
no maximum-paths
```

Description

Sets the maximum number of ECMP routes that RIP can support.

The **no** form of this command sets the maximum number of ECMP routes to the default value of 4.

Parameter	Description
<MAX-VALUE>	Sets the number of RIP ECMP routes. Range: 1-8.

Examples

Setting maximum number of RIP ECMP routes:

```
switch(config)# router rip 1
switch (config-rip-1)# maximum-paths 8
```

Setting maximum number of RIP ECMP routes to default:

```
switch(config)# router rip 1
switch (config-rip-1)# no maximum-paths
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

redistribute

```
redistribute {bgp | connected | ospf <PROCESS-ID> | static}
no redistribute {bgp | connected | ospf <PROCESS-ID> | static}
```

Description

Redistributes routes originating from other protocols into RIP.

The **no** form of this command disables redistribution of routes originating from other protocols into RIP.

Parameter	Description
`bgp`	Specifies BGP routes to redistribute into RIP.
`connected`	Specifies connected routes (directly attached subnet or host) to redistribute into RIP.
`ospf <PROCESS-ID>`	Specifies the OSPF route to redistribute into RIP. Range: <1-65535>
`static`	Specifies static route to redistribute into RIP.

Parameters **ospf** and **bgp** are not supported on 6100 switches.

Examples

Redistributing BGP routes into RIP:

```
switch(config)# router rip 1
switch(config-rip-1)# redistribute bgp
```

Disabling BGP routes that originate from other protocols and redistribute into RIP:

```
switch(config)# router rip 1
switch(config-rip-1)# no redistribute bgp
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

timers update

```
timers update <INTERVAL> timeout <DURATION> garbage-collection <PERIOD>
no timers
```

Description

Configures RIP timers with specific values.

The **no** form of this command sets all RIP timers to default values.

Parameter	Description
`timers update <INTERVAL>`	Specifies frequency at which RIP sends updates to all of its peers. Range: 1 to 2147484. Default: 30.
`timeout <DURATION>`	Specifies timeout duration from the point of the last refresh after a route is received from a peer timeout and is marked as expired. Range: 1 to 255. Default: 180.
`garbage-collection <PERIOD>`	Specifies amount of time route remains in routing table after route expiration. Range: 1 to 255. Default: 120.

Examples

Configuring RIP timers with specific values:

```
switch(config)# router rip 1
switch(config-rip-1)# timers update 40 timeout 200 garbage-collection 150
```

Configuring RIP timers with default values:

```
switch(config)# router rip 1
switch(config-rip-1)# no timers
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

RIPv2 clear commands

clear ip rip statistics

`clear ip rip [<PROCESS-ID>] statistics [all-vrfs | vrf <VRF-NAME>]`

Description

Clears RIP event statistics.

Parameter	Description
<PROCESS-ID>	Specifies RIP process ID. Range: 1-63
`all-vrfs`	Clears statistics for all VRFs.
`vrf`	Selects VRF to clear statistics for.
<VRF-NAME>	Specifies VRF name.

Examples

Clearing RIP event statistics:

```
switch# clear ip rip statistics
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

RIPv2 interface commands

enable

```
enable
no enable
```

Description

Enables RIP process for RIP enabled IP address configured on interface.

The **no** form of this command disables RIP process on interface.

Usage

- Default settings allow an interface to receive RIP packets.
- If an IP address is removed from an interface configured with RIP, all RIP configurations will be removed from the interface.

Examples

Enabling RIP process for RIP enabled IP address:

```
switch(config)# interface 1/1/1
switch(config-if)# ip rip 1 10.1.1.1
switch(config-if-rip)# enable
```

Disabling RIP process on interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ip rip 1 10.1.1.1
switch(config-if-rip)# no enable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if`	Administrators or local user group members with execution rights for this command.

disable

```
disable
no disable
```

Description

Disables RIP process for RIP enabled IP addresses configured on interface.

The **no** form of this command enables RIP process on interface.

Examples

Disabling RIP process for RIP enabled IP addresses configured on interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ip rip 1 10.1.1.1
switch(config-if-rip)# disable
```

Enabling RIP process on interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ip rip 1 10.1.1.1
switch(config-if-rip)# no disable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if`	Administrators or local user group members with execution rights for this command.

send disable

```
send disable
no send disable
```

Description

Disables an interface from sending RIP packets for a specific IP address.

The **no** form of this command enables interface for sending RIP packets for a specific IP address.

Usage

- Default settings allow an interface to send and receive RIP packets.
- If an IP address is removed from an interface configured with RIP, all RIP configurations will be removed from the interface.

Examples

Disabling interface from sending RIP packets for a specific IP address :

```
switch(config)# interface 1/1/1
switch(config-if)# ip rip 1 10.1.1.1
switch(config-if-rip)# send disable
```

Enabling interface to send RIP packets for a specific IP address:

```
switch(config)# interface 1/1/1
switch(config-if)# ip rip 1 10.1.1.1
switch(config-if-rip)# no send disable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if`	Administrators or local user group members with execution rights for this command.

receive disable

```
receive disable
no receive disable
```

Description

Disables interface from receiving RIP packets for a specific IP address.

The **no** form of this command enables interface for receiving RIP packets for a specific IP address.

Usage

- Default settings allow an interface to receive RIP packets.
- If an IP address is removed from an interface configured with RIP, all RIP configurations will be removed from the interface.

Examples

Disabling interface from receiving RIP packets for a specific IP address:

```
switch(config)# interface 1/1/1
switch(config-if)# ip rip 1 10.1.1.1
switch(config-if-rip)# receive disable
```

Enabling interface for receiving RIP packets for a specific IP address:

```
switch(config)# interface 1/1/1
switch(config-if)# ip rip 1 10.1.1.1
switch(config-if-rip)# no receive disable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

RIPv2 show commands

show capacities rip

```
show capacities rip
```

Description

Displays maximum number of RIP interfaces, routes and process.

Examples

Displaying maximum number of RIP interfaces, routes and process:

```
switch# show capacities rip

System Capacities: Filter RIP
Capacities Name                                                                  Value
-------------------------------------------------------------------------------------
Maximum number of RIP interfaces configurable in the system                      32
Maximum number of RIP processes supported across each VRF                        1
Maximum number of routes in RIP supported across all VRFs                        2540
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P	Operator (>) or Manager (#)	

Platforms	Command context	Authority
8360 9300 9300S 10000		

show capacities-status rip

```
show capacities-status rip
```

Description

Displays number of RIP interfaces, routes and process configured in the system.

Examples

Displaying number of RIP interfaces, routes and process:

```
switch# show capacities-status rip

System Capacities Status: Filter RIP
Capacities Name                                                 Value Maximum
---------------------------------------------------------------------------
Number of RIP interfaces configured in the system                   0      32
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ip rip

```
show ip rip [<PROCESS-ID>] [all-vrfs | vrf <VRF-NAME>]
```

Description

Displays general RIP configuration.

Parameter	Description
<PROCESS-ID>	Specifies RIP process ID. Range: 1-63.
`all vrfs`	Displays general RIP information for all VRFs.
`vrf`	Selects VRF to display general RIP information for.
<VRF-NAME>	Specifies VRF name.

Usage

- Parameters display general RIP information for a specific RIP process.
- Parameters display general RIP information for a specific or all VRFs.
- If a VRF is not mentioned, information for the default VRF is displayed.

Examples

Displaying general RIP configuration for all VRFs:

```
switch# show ip rip 34 all-vrfs
VRF : Default                        Process-ID : 34
------------------------------------------------------------------------
RIP Version             : RIPv2      Protocol Status : Enabled
Update Time             : 60  sec    Timeout Time    : 240 sec
Garbage Collection Time : 250 sec    ECMP            : 6
Distance                : 100        Redistribution  : static,
                                                       ospf 1

VRF : vrf_1                          Process-ID : 34
------------------------------------------------------------------------
RIP Version             : RIPv2      Protocol Status : Enabled
Update Time             : 30  sec    Timeout Time    : 180 sec
Garbage Collection Time : 120 sec    ECMP            : 4
Distance                : 120        Redistribution  : None
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300	Operator (>) or Manager (#)	

Platforms	Command context	Authority
9300S 10000		

show ip rip interface

```
show ip rip [<PROCESS-ID>] interface [<INTERFACE-NAME>] [brief] [all-vrfs | vrf <VRF-NAME>]
```

Description

Displays information about RIP enabled interfaces.

Parameter	Description
<PROCESS-ID>	Specifies RIP process ID. Range: 1-63.
<INTERFACE-NAME>	Specifies interface.
`brief`	Shows brief overview information for the RIP interface.
`all-vrfs`	Displays interface information for all VRFs.
`vrf`	Selects specific VRF.
<VRF-NAME>	Specifies VRF.

Usage

- Parameters display general RIP information for a specific RIP process.
- If a VRF is not mentioned, information for the default VRF is displayed.

Examples

```
switch# show ip rip interface
Interface 1/1/1 is up, IP Address is 10.10.10.1/24
------------------------------------------------------------------------
VRF             : Default            Process-ID     : 1
Status          : Oper Up            Mode           : Send and Receive
MTU             : 500                Version        : RIPv2
Poision Reverse : Enabled

Interface 1/1/2 is up, IP Address is 20.10.10.1/24
------------------------------------------------------------------------
VRF             : Default            Process-ID     : 1
Status          : Admin Down         Mode           : Receive
MTU             : 500                Version        : RIPv2
Poision Reverse : Enabled

Interface 1/1/3 is up, IP Address is 30.10.10.1/24
------------------------------------------------------------------------
VRF             : Default            Process-ID     : 1
Status          : Admin Down         Mode           : Send
MTU             : 500                Version        : RIPv2
Poision Reverse : Enabled

switch# show ip rip interface brief
```

```
VRF : default   Process-ID : 1
-------------------------------

Total Number of Interfaces: 2

Interface     IP-Address/Mask    Status    MTU
-------------------------------------------------
1/1/1         10.10.10.1/24       up       500
1/1/2         20.10.10.1/24       up       500
1/1/3         30.10.10.1/24       up       500
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ip rip neighbors

`show ip rip [<PROCESS-ID>] neighbors [<IP-ADDRESS>] [all-vrfs | vrf <VRF-NAME>]`

Description

Displays information about RIP neighbors.

Parameter	Description
<PROCESS-ID>	Specifies RIP process ID. Range: 1-63.
<IP-ADDRESS>	Specifies IP address of a specific neighbor to display information on.
`all-vrfs`	Displays neighbor information for all VRFs.
`vrf`	Selects VRF to display neighbor information.
<VRF-NAME>	Specifies VRF name.

Usage

- Parameters display RIP neighbor information for a specific RIP process.
- Parameters display RIP neighbor information for a specific neighbor.
- If a VRF is not mentioned, information for the default VRF is displayed.

Examples

Displaying RIP neighbor information for all VRFs:

```
switch# show ip rip neighbors all-vrfs
VRF : default            Process-ID : 1
----------------------------------------

Total Number of Neighbors: 1

Peer-Address    Type    Last-Update  Rcvd-Bad-Pkts    Rcvd-Bad-Routes
-----------------------------------------------------------------
1.1.1.2         RIPv2   0:0:7        4                5
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ip rip routes

`show ip rip [<PROCESS-ID>] routes [<PREFIX/LENGTH>] [all-vrfs | vrf <VRF-NAME>]`

Description

Displays RIP routing table for a specific RIP process.

Parameter	Description
<PROCESS-ID>	Specifies RIP process ID to display information for a specific RIP process. Range: 1-63.
<PREFIX/LENGTH>	Specifies the network prefix.

Parameter	Description
`all-vrfs`	Displays RIP routing information for all VRFs.
`vrf`	Selects VRF to display RIP routing information.
<VRF-NAME>	Specifies VRF name.

Usage

- *<PREFIX/LENGTH>* is an optional parameter that displays RIP routing table information for a specific subnet.
- If a VRF is not mentioned, information for the default VRF is displayed.

Examples

Displaying RIP routing table for all VRFs:

```
switch# show ip rip routes all-vrfs
VRF : default              Process-ID : 1
---------------------------------------
Total Number of Routes : 6

Prefix              Metric     Interface      Nexthop
-------------------------------------------------------
10.1.0.0/16           2         1/1/1         30.1.1.2
20.1.2.0/24           3         1/1/1         30.1.1.2
30.1.1.0/24           1         1/1/1

VRF : vrf_1                Process-ID : 34
---------------------------------------

Prefix               Metric     Interface      Nexthop
-------------------------------------------------------
20.1.0.0/16           10         1/1/2         50.1.1.2
40.1.2.0/24           14         1/1/2         50.1.1.2
50.1.1.0/24            1         1/1/2
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360	Operator (>) or Manager (#)	

Platforms	Command context	Authority
9300 9300S 10000		

show ip rip statistics

```
show ip rip [<PROCESS-ID>] statistics [all-vrfs | vrf <VRF-NAME>]
```

Description

Displays RIP statistics.

Parameter	Description
<PROCESS-ID>	Specifies RIP process ID. Range: 1-63.
`all-vrfs`	Displays statistics information for all VRFs.
`vrf`	Selects VRF to display RIP statistics information for.
<VRF-NAME>	Specifies VRF name.

Usage

- Parameters can display information for all VRFs or a specific VRF.
- If a VRF is not mentioned, information for the default VRF is displayed.

Examples

Displaying RIP statistics for all VRFs:

```
switch# show ip rip statistics all-vrfs
VRF : default  Process-ID : 1
-------------------------------------

Global Route Changes : 50
Global Queries       : 2
Last Cleared         : 0h 30m 28s ago

VRF : vrf_1    Process-ID : 34
-------------------------------------

Global Route Changes : 20
Global Queries       : 0
Last Cleared         : 0h 30m 28s ago
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ip rip statistics interface

```
show ip rip [<PROCESS-ID>] statistics interface [<INTERFACE-NAME>] [all-vrfs | vrf <VRF-NAME>]
```

Description

Displays RIP statistics for RIP enabled interfaces.

Parameter	Description
<PROCESS-ID>	Specifies RIP process ID. Range: 1-63.
<INTERFACE-NAME>	Specifies name of interface.
`all-vrfs`	Displays RIP interface statistics for all VRFs.
`vrf`	Selects VRF to display RIP interface statistics.
<VRF-NAME>	Specifies VRF name.

Usage

- Parameters can display information for all VRFs or a specific VRF.
- If a VRF is not mentioned, information for the default VRF is displayed.

Examples

Displaying RIP statistics for a RIP enabled interface:

```
switch# show ip rip statistics interface 1/1/1
VRF : default  Process-ID : 1      interface 1/1/1
-------------------------------------------------

IP-Address    Trigger-Updates    Rcvd-Bad-Packets    Rcvd-Bad-Routes
-------------------------------------------------------------------------
10.1.1.1      15                 3                   4

Last Cleared : 0h 30m 28s ago
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show running-config

```
show running config
```

Description

Displays all running configurations for all protocols including RIP.

Examples

Displaying all running configurations for all protocols including RIP:

```
switch# show running-config
Current configuration:
!
!Version Halon 0.1.0 (Build: genericx86-64-Halon-0.1.0-master-20170309054955-dev)
!Schema version 0.1.8
lldp enable
timezone set utc
vrf blue
vrf green
vrf red
led base-loc_fdc on
led base-loc on
led base-hlth_fdc fast_blink
led base-pwr_fdc on
!
!
!
!
!
!
aaa authentication login default local
aaa authorization commands default none
!
!
!
!
router ospf 1 vrf red
router rip 1
```

```
    maximum-paths 5
    distance 1
router rip 1 vrf red
    default-information originate always
    maximum-paths 7
    distance 5
    redistribute ospf 1
    timers update 40 timeout 200 garbage-collection 120
vlan 1
    no shutdown
interface lag 44
    no shutdown
    ip address 33.1.1.1/24
    ip rip 1 33.1.1.1
        send disable
interface 1/1/1
    no shutdown
    ip address 33.44.1.1/24
    ip address 44.44.1.1/24 secondary
    ip rip 1 33.44.1.1
        send disable
    ip rip 1 44.44.1.1
        send disable
interface 1/1/2
interface loopback 2
    ip address 55.55.55.55/32
    ip rip 1 55.55.55.55
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

RIPng (IPv6) commands

Configuration commands

router ripng

`router ripng <PROCESS-ID> [vrf <VRF-NAME>]`

```
no router ripng <PROCESS-ID> [vrf <VRF-NAME>]
```

Description

Creates RIPng process if not already created and enters the **router ripng <PROCESS-ID>** context for the VRF mentioned. If no VRF is mentioned, a default is used. Only one RIPng process is allowed per VRF.

The **no** form of this command deletes the RIPng instance for the VRF. If no VRF is mentioned the default is deleted.

Parameter	Description
<PROCESS-ID>	Specifies name of the RIPng process ID. Range: <1-63>
vrf	Sets VRF name for RIPng process.
<VRF-NAME>	VRF name for VRF.

Examples

Creating RIPng process and naming the VRF:

```
switch(config)# router ripng 2 vrf red
```

Deleting RIPng process:

```
switch(config)# no router ripng 2 vrf red
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config	Administrators or local user group members with execution rights for this command.

Interface commands

ipv6 ripng

```
ipv6 ripng <PROCESS-ID>
no ipv6 ripng <PROCESS-ID>
```

Description

Enables RIPng process on interface and creates a new context.

The **no** form of this command deletes RIPng process on interface.

Parameter	Description
<PROCESS-ID>	Specifies RIPng process ID. Range: 1-63.

Examples

Enabling RIPng process on an interface:

```
switch(config)# interface 1/1/1
switch (config-if)# ipv6 ripng 1
switch (config-if-ripng)#
```

Deleting RIPng process on an interface:

```
switch(config)# interface 1/1/1
switch (config-if)# no ipv6 ripng 1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

Routing commands

enable

```
enable
no enable
```

Description

Enables RIPng process if disabled. By default RIPng process is enabled.

The **no** form of this command disables the RIPng process.

Examples

Enabling RIPng process when disabled:

```
switch(config)# router ripng 1
switch(config-ripng-1)# enable
```

Disabling RIPng process when enabled:

```
switch(config)# router ripng 1
switch(config-ripng-1)# no enable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

disable

```
disable
no disable
```

Description

Disables RIPng process.

The **no** form of this command enables the RIPng process.

Examples

Disabling RIPng process:

```
switch(config)# router ripng 1
switch(config-ripng-1)# disable
```

Enabling RIPng process:

```
switch(config)# router ripng 1
switch(config-ripng-1)# no disable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

distance

```
distance <DISTANCE>
no distance
```

Description

Configures administrative distance for RIPng. Administrative distance is used as criteria to select the best route when multiple protocols have the same route.

The **no** form of this command sets the RIPng administrative distance to the default. Default: 120.

Parameter	Description
<DISTANCE>	Specifies RIPng administrative distance. Range: 1 to 255.

Examples

Configuring administrative distance for RIPng:

```
switch(config)# router ripng 1
switch(config-ripng-1)# distance 100
```

Setting administrative distance for RIPng to default values:

```
switch(config)# router ripng 1
switch(config-ripng-1)# no distance
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config	Administrators or local user group members with execution rights for this command.

maximum-paths

```
maximum-paths <MAX-VALUE>
no maximum-paths
```

Description

Sets the maximum number of ECMP routes that RIPng can support.

The **no** form of this command sets the maximum number of ECMP routes to the default value of 4.

Parameter	Description
<MAX-VALUE>	Sets the number of RIPng ECMP routes. Range: 1-8.

Examples

Setting maximum number of RIPng ECMP routes:

```
switch(config)# router ripng 1
switch (config-ripng-1)# maximum-paths 8
```

Setting maximum number of RIPng ECMP routes to default:

```
switch(config)# router ripng 1
switch (config-ripng-1)# no maximum-paths
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

redistribute

```
redistribute {bgp | connected | ospfv3 <PROCESS-ID> | static}
no redistribute {bgp | connected | ospfv3 <PROCESS-ID> | static}
```

Description

Redistributes routes originating from other protocols into RIPng.

The **no** form of this command disables redistribution of routes originating from other protocols into RIPng.

Parameter	Description
`bgp`	Specifies BGP routes to redistribute into RIPng.
`connected`	Specifies connected routes (directly attached subnet or host) to redistribute into RIPng.
`ospfv3 <PROCESS-ID>`	Specifies the OSPFv3 route to redistribute into RIPng. Range: <1-65535>
`static`	Specifies static route to redistribute into RIPng.

Parameters **ospfv3** and **bgp** are not supported on 6100 switches.

Examples

Redistributing BGP routes into RIPng:

```
switch(config)# router ripng 1
switch(config-ripng-1)# redistribute bgp
```

Disabling BGP routes that originate from other protocols and redistribute into RIPng:

```
switch(config)# router ripng 1
switch(config-ripng-1)# no redistribute bgp
```

Command History

Release	Modification
10.14	Supported process ID range expanded from 1-63 to 1-65535.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

timers update

`timers update <INTERVAL> timeout <DURATION> garbage-collection <PERIOD>`
`no timers`

Description

Configures RIPng timers with specific values.

The **no** form of this command sets all RIPng timers to default values.

Parameter	Description
`timers update <INTERVAL>`	Specifies frequency at which RIPng sends updates to all of its peers. Range: 1 to 2147484. Default: 30.
`timeout <DURATION>`	Specifies timeout duration from the point of the last refresh after a route is received from a peer timeout and is marked as expired. Range: 1 to 255. Default: 180.

Parameter	Description
`garbage-collection <PERIOD>`	Specifies amount of time route remains in routing table after route expiration. Range: 1 to 255. Default: 120.

Examples

Configuring RIPng timers with specific values:

```
switch(config)# router ripng 1
switch(config-ripng-1)# timers update 40 timeout 200 garbage-collection 150
```

Configuring RIPng timers with default values:

```
switch(config)# router ripng 1
switch(config-ripng-1)# no timers
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config`	Administrators or local user group members with execution rights for this command.

RIPng clear commands

clear ipv6 ripng statistics

`clear ipv6 ripng [<PROCESS-ID>] statistics [all-vrfs | vrf <VRF-NAME>]`

Description

Clears RIPng event statistics.

Parameter	Description
<PROCESS-ID>	Specifies RIPng process ID. Range: 1-63
`all-vrfs`	Clears statistics for all VRFs.
`vrf`	Selects VRF to clear statistics for.
<VRF-NAME>	Specifies VRF name.

Examples

Clearing RIPng event statistics:

```
switch# clear ipv6 ripng statistics
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

RIPng interface commands

enable

```
enable
no enable
```

Description

Enables RIPng process on interface.

The **no** form of this command disables RIPng process on interface.

Examples

Enabling RIPng process on interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ipv6 ripng 1
switch(config-if-ripng)# enable
```

Disabling RIPng process on interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ipv6 ripng 1
switch(config-if-ripng)# no enable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

disable

```
disable
no disable
```

Description

Disables RIPng process on interface.

The **no** form of this command enables RIPng process on interface.

Examples

Disabling RIP process for all RIP enabled IP addresses on interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ipv6 ripng 1
switch(config-if-ripng)# disable
```

Enabling RIP process for all RIP enabled IP addresses on interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ipv6 ripng 1
switch(config-if-ripng)# no disable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	config-if	Administrators or local user group members with execution rights for this command.

send disable

```
send disable
no send disable
```

Description

Disables interface from sending RIPng packets. An interface can send RIPng packets by default.

The **no** form of this command enables interface to send RIPng packets, if disabled.

Examples

Disabling interface from sending RIPng packets:

```
switch(config)# interface 1/1/1
switch (config-if)# ipv6 ripng 1
switch (config-if-ripng)# send disable
```

Enabling interface to send RIPng packets:

```
switch(config)# interface 1/1/1
switch (config-if)# ipv6 ripng 1
switch (config-if-ripng)# no send disable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if`	Administrators or local user group members with execution rights for this command.

receive disable

```
receive disable
no receive disable
```

Description

Disables interface from receiving RIPng packets for all enabled IP addresses. An interface can receive RIPng packets by default.

The **no** form of this command enables interface to receive RIPng packets, if disabled.

Examples

Disabling interface from receiving RIPng packets:

```
switch(config)# interface 1/1/1
switch(config-if)# ipv6 ripng 1
switch (config-if-ripng)# receive disable
```

Enabling interface to receive RIPng packets when disabled:

```
switch(config)# interface 1/1/1
switch(config-if)# ipv6 ripng 1
switch (config-if-ripng)# no receive disable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	`config-if`	Administrators or local user group members with execution rights for this command.

RIPng show commands

show capacities ripng

```
show capacities ripng
```

Description

Displays the maximum number of RIPng interfaces, routes and process.

Examples

Displaying maximum number of RIPng interfaces, routes and process:

```
switch# show capacities ripng

System Capacities: Filter RIPng
Capacities Name                                                              Value
-----------------------------------------------------------------------------------
Maximum number of RIPng interfaces configurable in the system                32
Maximum number of RIPng processes supported across each VRF                  1
Maximum number of routes in RIPng supported across all VRFs                  2540
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P	Operator (>) or Manager (#)	

Platforms	Command context	Authority
8360 9300 9300S 10000		

show capacities-status ripng

```
show capacities-status ripng
```

Description

Displays number of RIPng interfaces, routes and process configured in the system.

Examples

Displaying number of RIPng interfaces, routes and process:

```
switch# show capacities-status ripng

System Capacities Status: Filter RIPng
Capacities Name
          Value  Maximum
-------------------------------------------------------------------------------
-------------------------
Number of RIPng interfaces configured in the system
              0       32
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ipv6 ripng

```
show ipv6 ripng [<PROCESS-ID>] [all-vrfs | vrf <VRF-NAME>]
```

Description

Displays general RIPng configuration.

Parameter	Description
<PROCESS-ID>	Specifies RIPng process ID. Range: 1-63.
`all vrfs`	Displays general RIPng information for all VRFs.
`vrf`	Selects VRF to display general RIPng information for.
<VRF-NAME>	Specifies VRF name.

Usage

- Parameters display general RIPng information for a specific RIPng process.
- Parameters display general RIPng information for a specific or all VRFs.
- If a VRF is not mentioned, information for the default VRF is displayed.

Examples

Displaying general RIPng configuration for all VRFs:

```
switch# show ipv6 ripng 34 all-vrfs
VRF : Default                         Process-ID : 34
------------------------------------------------------------------------
Protocol Status         : Enabled   ECMP            : 6
Update Time             : 60  sec   Timeout Time    : 240 sec
Garbage Collection Time : 250 sec   Distance        : 100
Redistribution          : static,
                          ospfv3 1

VRF : vrf_1                           Process-ID : 34
------------------------------------------------------------------------
Protocol Status         : Enabled   ECMP            : 4
Update Time             : 30  sec   Timeout Time    : 180 sec
Garbage Collection Time : 120 sec   Distance        : 120
Redistribution          : None
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300	Operator (>) or Manager (#)	

Platforms	Command context	Authority
9300S 10000		

show ipv6 ripng interface

```
show ipv6 ripng [<PROCESS-ID>] interface [<INTERFACE-NAME>] [brief] [all-vrfs | vrf <VRF-NAME>]
```

Description

Displays information about RIPng enabled interfaces.

Parameter	Description
<PROCESS-ID>	Specifies RIPng process ID. Range: 1-63.
<INTERFACE-NAME>	Specifies interface.
`brief`	Shows brief overview information for RIPng interface.
`all-vrfs`	Displays interface information for all VRFs.
`vrf`	Selects specific VRF.
<VRF-NAME>	Specifies VRF.

Usage

- Parameters display general RIPng information for a specific RIPng process.
- Parameters display general RIPng information for a specific or all VRFs.
- If a VRF is not mentioned, information for the default VRF is displayed.

Examples

```
switch# show ipv6 ripng interface
Interface 1/1/1 is up, IPv6 Address is fe80::7272:cfff:fe70:67a
------------------------------------------------------------------------
VRF             : Default            Process-ID      : 1
Status          : Oper Up            Mode            : Send and Receive
MTU             : 500                Poision Reverse : Enabled

Interface 1/1/2 is up, IPv6 Address is fe80::7272:cfff:fe70:67a
------------------------------------------------------------------------
VRF             : Default            Process-ID      : 1
Status          : Admin Down         Mode            : Receive
MTU             : 500                Poision Reverse : Enabled

Interface 1/1/3 is up, IPv6 Address is fe80::7272:cfff:fe70:67a
------------------------------------------------------------------------
VRF             : Default            Process-ID      : 1
Status          : Admin Down         Mode            : Send
MTU             : 500                Poision Reverse : Enabled

switch# show ipv6 ripng interface brief
VRF : default   Process-ID : 1
```

```
-------------------------------

Total Number of Interfaces: 2

Interface     IPv6-Address                    Status    MTU
-------------------------------------------------------------
1/1/1         fe80::7272:cfff:fe70:67a        up        500
1/1/2         fe80::7272:cfff:fe71:67a        up        500
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ipv6 ripng neighbors

`show ipv6 ripng [<PROCESS-ID>] neighbors [<LINK-LOCAL-ADDRESS>] [all-vrfs | vrf <VRF-NAME>]`

Description

Displays information about RIPng neighbors.

Parameter	Description
<PROCESS-ID>	Specifies RIPng process ID. Range: 1-63.
`neighbors`	Specifies neighbor IP address.
<LINK-LOCAL-ADDRESS>	Specifies link-local address.
`all-vrfs`	Displays neighbor information for all VRFs.
`vrf`	Selects VRF to display neighbor information.
<VRF-NAME>	Specifies VRF name.

Usage

- Parameters display RIPng neighbor information for a specific RIPng process.
- Parameters display RIPng neighbor information for a specific neighbor.
- Parameters display general RIPng information for a specific or all VRFs.
- If a VRF is not mentioned, information for the default VRF is displayed.

Examples

Displaying RIPng neighbor information for all VRFs:

```
switch# show ipv6 ripng neighbors all-vrfs
VRF : default            Process-ID : 1
----------------------------------------

Total Number of Neighbors: 1

Peer-Address     Type    Last-Update  Rcvd-Bad-Pkts    Rcvd-Bad-Routes
-------------------------------------------------------------------------
fe80::7272:cfff:fe70:86ae
                 RIPng   0:0:7        4                5
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ipv6 ripng routes

`show ipv6 ripng [<PROCESS-ID>] routes [<PREFIX/LENGTH>] [all-vrfs | vrf <VRF-NAME>]`

Description

Displays RIPng routing table for a specific RIPng process.

Parameter	Description
<PROCESS-ID>	Specifies RIPng process ID to display information for a specific RIPng process. Range: 1-63.

Parameter	Description
<PREFIX/LENGTH>	Specifies the network prefix.
`all-vrfs`	Displays RIPng routing information for all VRFs.
`vrf`	Selects VRF to display RIPng routing information.
<VRF-NAME>	Specifies VRF name.

Usage

- *<PREFIX/LENGTH>* is an optional parameter that displays RIPng routing table information for a specific subnet.
- If a VRF is not mentioned, information for the default VRF is displayed.

Examples

Displaying RIPng routing table for all VRFs:

```
switch# show ipv6 ripng routes all-vrfs
VRF : default              Process-ID : 1
-----------------------------------------

Prefix              Metric     Interface      Nexthop
-------------------------------------------------------------------------
2001:DB8:10::/64      2         1/1/1         FE80::2E0:E6FF:FE1B:8242
2002:DB8:10::/64      3         1/1/1         FE80::2E0:E6FF:FE1B:8242
2003:DB8:10::/64      1         1/1/1

VRF : vrf_1                Process-ID : 34
-----------------------------------------

Prefix               Metric     Interface      Nexthop
-------------------------------------------------------------------------
3001:DB8:10::/64      10         1/1/2         FE80::2E0:E6FF:FE1B:8232
3002:DB8:10::/64      14         1/1/2         FE80::2E0:E6FF:FE1B:8232
3003:DB8:10::/64       1         1/1/2
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P	Operator (>) or Manager (#)	

Platforms	Command context	Authority
8360 9300 9300S 10000		

show ipv6 ripng statistics

```
show ipv6 ripng [<PROCESS-ID>] statistics [all-vrfs | vrf <VRF-NAME>]
```

Description

Displays RIPng statistics.

Parameter	Description
<PROCESS-ID>	Specifies RIPng process ID. Range: 1-63.
all-vrfs	Displays statistics information for all VRFs.
vrf	Selects VRF and displays RIPng statistics for it.
<VRF-NAME>	Specifies VRF name.

Usage

- Parameters can display information for all VRFs or a specific VRF.
- If a VRF is not mentioned, information for the default VRF is displayed.

Examples

Displaying RIPng statistics for all VRFs:

```
switch# show ipv6 ripng statistics all-vrfs
VRF : default  Process-ID : 1
-------------------------------------

Global Route Changes : 50
Global Queries       : 2
Last Cleared         : 0h 30m 28s ago

VRF : vrf_1    Process-ID : 34
-------------------------------------

Global Route Changes : 20
Global Queries       : 0
Last Cleared         : 0h 30m 28s ago
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show ipv6 ripng statistics interface

```
show ipv6 ripng [<PROCESS-ID>] statistics interface [<INTERFACE-NAME>] [all-vrfs | vrf
<VRF-NAME>]
```

Description

Displays RIPng statistics for RIPng enabled interfaces.

Parameter	Description
<PROCESS-ID>	Specifies RIPng process ID. Range: 1-63.
<INTERFACE-NAME>	Specifies name of interface.
`all-vrfs`	Displays RIPng interface statistics for all VRFs.
`vrf`	Selects VRF to display RIPng interface statistics.
<VRF-NAME>	Specifies VRF name.

Usage

- Parameters can display information for all VRFs or a specific VRF.
- If a VRF is not mentioned, information for the default VRF is displayed.

Examples

Displaying RIPng statistics for a RIPng enabled interface:

```
switch# show ipv6 ripng statistics interface 1/1/1
VRF : default  Process-ID : 1      interface 1/1/1
-------------------------------------------------

IPv6-Address          Trigger-Updates    Rcvd-Bad-Packets    Rcvd-Bad-Routes
-------------------------------------------------------------------------------
fe80::7272:cfff:fe70:86ae
                      15                 3                   4

Last Cleared: 0h 30m 28s ago
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

show running-config

```
show running config
```

Description

Displays all running configurations for all protocols including RIPng.

Examples

Displaying all running configurations for all protocols including RIPng:

```
switch# show running-config
Current configuration:
!
!Version Halon 0.1.0 (Build: genericx86-64-Halon-0.1.0-master-20170309054955-dev)
!Schema version 0.1.8
lldp enable
timezone set utc
vrf green
vrf red
led base-loc_fdc on
led base-loc on
led base-hlth_fdc fast_blink
led base-pwr_fdc on
!
!
!
!
!
!
aaa authentication login default local
aaa authorization commands default none
!
!
!
!
router ospfv3 1
router ripng 1
    maximum-paths 5
```

```
    distance 1
    redistribute ospfv3 1
    timers update 40 timeout 200 garbage-collection 150
router ripng 1 vrf red
    default-information originate always
    maximum-paths 7
    distance 5
vlan 1
    no shutdown
interface lag 44
    no shutdown
    ipv6 address link-local
    ipv6 ripng 1
        send disable
interface 1
    no shutdown
    ipv6 address link-local
    ipv6 ripng 1
        receive disable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A 6400 8100 8325 8325H 8325P 8360 9300 9300S 10000	Operator (>) or Manager (#)	

Chapter 18

Support and Other Resources

Accessing HPE Aruba Networking Support

HPE Aruba Networking Support Services	https://www.hpe.com/us/en/networking/hpe-aruba-networking-support-services.html
AOS-CX Switch Software Documentation Portal	https://arubanetworking.hpe.com/techdocs/AOS-CX/help_portal/Content/home.htm
HPE Aruba Networking Support Portal	https://networkingsupport.hpe.com/home
North America telephone	1-800-943-4526 (US & Canada Toll-Free Number) +1-650-750-0350 (Backup—Toll Number)
International telephone	https://www.hpe.com/psnow/doc/a50011948enw

Be sure to collect the following information before contacting Support:

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Other useful sites

Other websites that can be used to find information:

HPE Aruba Networking Developer Hub	https://developer.arubanetworks.com/hpe-aruba-networking-aoscx/docs/about
Airheads social forums and Knowledge Base	https://community.arubanetworks.com/
AOS-CX Software Technical Update channel on YouTube.	Videos on new features introduced in this release: https://www.youtube.com/playlist?list=PLsYGHuNuBZcbWPEjjHuVMqP-Q_UL3CskS

HPE Aruba Networking Hardware Documentation and Translations Portal	https://arubanetworking.hpe.com/techdocs/hardware/DocumentationPortal/Content/home.htm
HPE Aruba Networking software	https://networkingsupport.hpe.com/downloads
Software licensing and Feature Packs	https://licensemanagement.hpe.com/
End-of-Life information	https://networkingsupport.hpe.com/end-of-life

Accessing Updates

You can access updates from the HPE Aruba Networking Support Portal at https://networkingsupport.hpe.com.

Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.

To subscribe to eNewsletters and alerts:

https://networkingsupport.hpe./notifications/subscriptions (requires an active HPE Aruba Networking Support Portal account to manage subscriptions). Security notices are viewable without an HPE Aruba Networking Support Portal account.

Warranty Information

To view warranty information for your product, go to https://www.arubanetworks.com/support-services/product-warranties/.

Regulatory Information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at https://www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional regulatory information

HPE Aruba Networking is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements, environmental data (company programs, product recycling, energy efficiency), and safety information and compliance data, (RoHS and WEEE). For more information, see https://www.arubanetworks.com/company/about-us/environmental-citizenship/.

Documentation Feedback

HPE Aruba Networking is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback

(docsfeedback-switching@hpe.com). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.